A Unified Cognitive Model of Visual Filling-In Based on an Emergic Network Architecture

by

David Pierre Leibovitz

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Cognitive Science

Carleton University
Ottawa, Ontario

Abstract

The Emergic Cognitive Model (ECM) is a unified computational model of visual filling-in based on the Emergic Network architecture. The Emergic Network was designed to help realize systems undergoing continuous change. In this thesis, eight different filling-in phenomena are demonstrated under a regime of continuous eye movement (and under static eye conditions as well).

ECM indirectly demonstrates the power of unification inherent with Emergic Networks when cognition is decomposed according to finer-grained functions supporting change. These can interact to raise additional emergent behaviours via cognitive re-use, hence the *Emergic* prefix throughout. Nevertheless, the model is robust and parameter free. Differential re-use occurs in the nature of model interaction with a particular testing paradigm.

ECM has a novel decomposition due to the requirements of handling motion and of supporting unified modelling via finer functional grains. The breadth of phenomenal behaviour covered is largely to lend credence to our novel decomposition.

The Emergic Network architecture is a hybrid between classical connectionism and classical computationalism that facilitates the construction of unified cognitive models. It helps *cutting up* of functionalism into finer-grains distributed over space (by harnessing massive recurrence) and over time (by harnessing continuous change), yet simplifies by using standard computer code to focus on the interaction of information flows. Thus

while the structure of the network looks *neurocentric*, the dynamics are best understood in *flowcentric* terms. Surprisingly, dynamic system analysis (as usually understood) is *not* involved. An Emergic Network is engineered much like straightforward software or hardware systems that deal with continuously varying inputs. Ultimately, this thesis addresses the problem of reduction and induction over complex systems, and the Emergic Network architecture is merely a tool to assist in this epistemic endeavour.

ECM is strictly a *sensory* model and apart from *perception*, yet it is informed by phenomenology. It addresses the attribution problem of how much of a phenomenon is best explained at a sensory level of analysis, rather than at a perceptual one. As the causal information flows are stable under eye movement, we hypothesize that they are the locus of consciousness, howsoever it is ultimately realized.

Acknowledgements

Thanks to Tony White who helped me begin this path; to Robert West who helped me finish; to Jo, Nikita, Keara, PaddyPaws and Seiji who continue to support me on this never-ending journey; to innumerable individuals at Carleton who inspired me; to Mike who entertained me; and to Heraclitus who changed me.

Table of Contents (Top-Level)

Table of Contents

List of Illustrations

List of Appendices

1   Chapter: Introduction

This thesis presents a model of visual filling-in under conditions that include continuous

eye motion whereby the different filling-in phenomena are produced as emergent

properties of the interactions among the underlying functions concerned with *maintaining*

information coherence and *handling missing data*. These *emergic* functions interact to

ensure that each visual processing layer compensates for eye motion by behaving as if it

was a "shift register". However, unlike traditional *feedforward* shifting circuits that

stabilize visual content onto a *final* layer of spatiotopic neurons (e.g., C. H. Anderson &

Van Essen, 1987), ours maintains coherence in the "flow of information" between *every*

inter-layer connection including *recurrent* ones. This inadvertently led to behaviour that

matches the neurobiology of receptive field remapping (Merriam, Genovese, & Colby,

2007). Each of our layers can be thought of as maintaining a separate memory or "model

of the world" which is updated by coherent information feeding upwards from lower

layers, or recurring downwards from higher ones.

To facilitate a decompositional scheme that focuses on interaction, the emergic network

modeling system was created. It is a species of artificial neural networks that deals

explicitly with recurrent interactions within each unit.

1.1   Emergic network modeling system

The emergic network modeling system is designed to explore how lower level functions

are combined to produce multiple instances of higher level functionality. Emergic

networks are non-symbolic and in some ways similar to spiking neural models. In

particular, the emergic units that make up the networks are similar to clusters of neurons
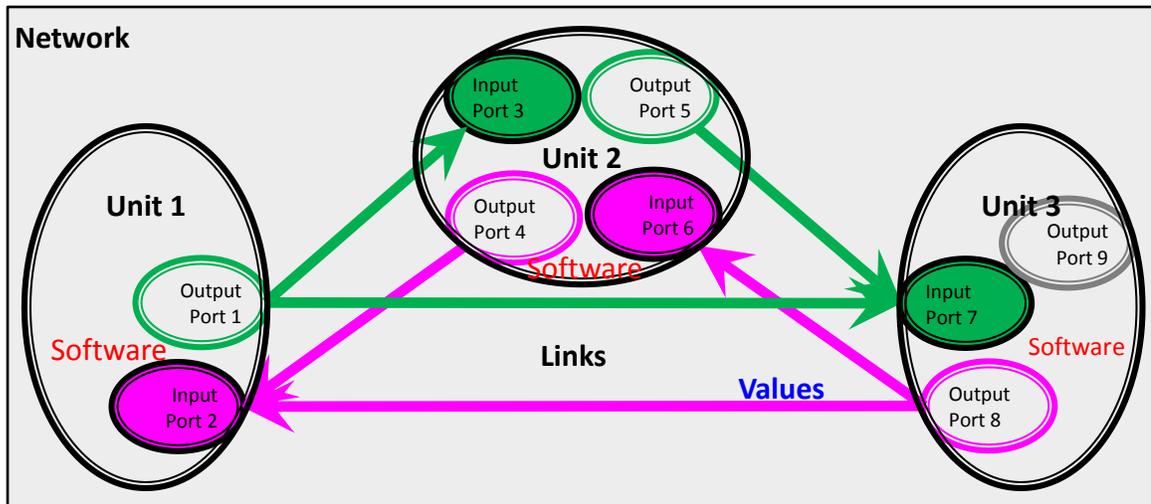
Illustration 1: Emergic network example

that perform a specific function, and the connections between the emergic units are functionally similar to the neural connections between clusters. Also, similar to actual neural systems, emergic networks process information in a continuous manner. However, the point of emergic networks is to model emergence based on higher level functional interactions rather than on detailed neural interactions.

An *emergic network* (e.g., Illustration 1) consists of a connected set of *emergic units*, each forming a micro-level portion of functional computation and behaviour. An emergic unit computes a function that can be represented mathematically or by computer code. Emergic units have input and output *ports* that connect them to other emergic units through *links*. Links transport *values* between the units. The values can take any mathematical form (e.g., numbers, vectors, and statistics). Links are unidirectional and have a scaling factor that is set to 1 by default. Input and output ports can be connected to multiple links. By default, link values are summed at input ports while an output port will duplicate its value to all destinations. The emergic units are distinguished from traditional connectionist units (also see Illustration 2) in having two or more outputs or flows. Thus

2

the focus of decomposition is in computing the interaction among flows. It is this focus on massive interactivity (and recurrence) –which is known to lead to emergent effects – that provides the *emergic* prefix throughout the thesis.

The emergic network architecture is synchronous, with links having a minimal delay of one tick. That is, the delivery of values through the links is clocked so that all values arrive at the same time. The values flowing around the network are intended to represent small changes, i.e., to approximate a physical system of continuous change and interaction (Rumelhart, McClelland, & Group, 1986). Emergic networks must model asynchronous behaviour by setting time delays (counted in ticks) small enough for computing the effective functions of emergic units in an incremental fashion.

The code for building Emergic Networks can be downloaded from:

```
http://emergic.upwize.com/?page_id=6
```

Details about the emergic modeling system are available in Appendix F. A realistic model of the human eye (see Illustration 9, p15) was also used and is described within Appendix G. The virtual eye was capable of looking at anything presented on a computer screen. This model/environment interaction system is described in Appendix H.

The thesis describes how the model produces the phenomena of colour homogeneity, trans-saccadic integration, temporal filling-in after a brief presentation and spatial filling-in through the optical blind-spot. Other related phenomena that the model can produce, but could not be included in the body of the thesis due to space restrictions, are

anorthoscopic perception, filling-in of the foveal blue scotoma, temporal filling-in thru blinks and imagination. These results are presented in Appendixes A, B, C and D respectively. The model is summarized next with additional details within the body of the thesis and with full details in Appendix G.

A significant product of this research is the concept of *flowcentric* analysis (Illustration 2). Units are considered to be stateless, while visual content is maintained within the flow of values. This maintenance leads to the emergence of memory with novel properties
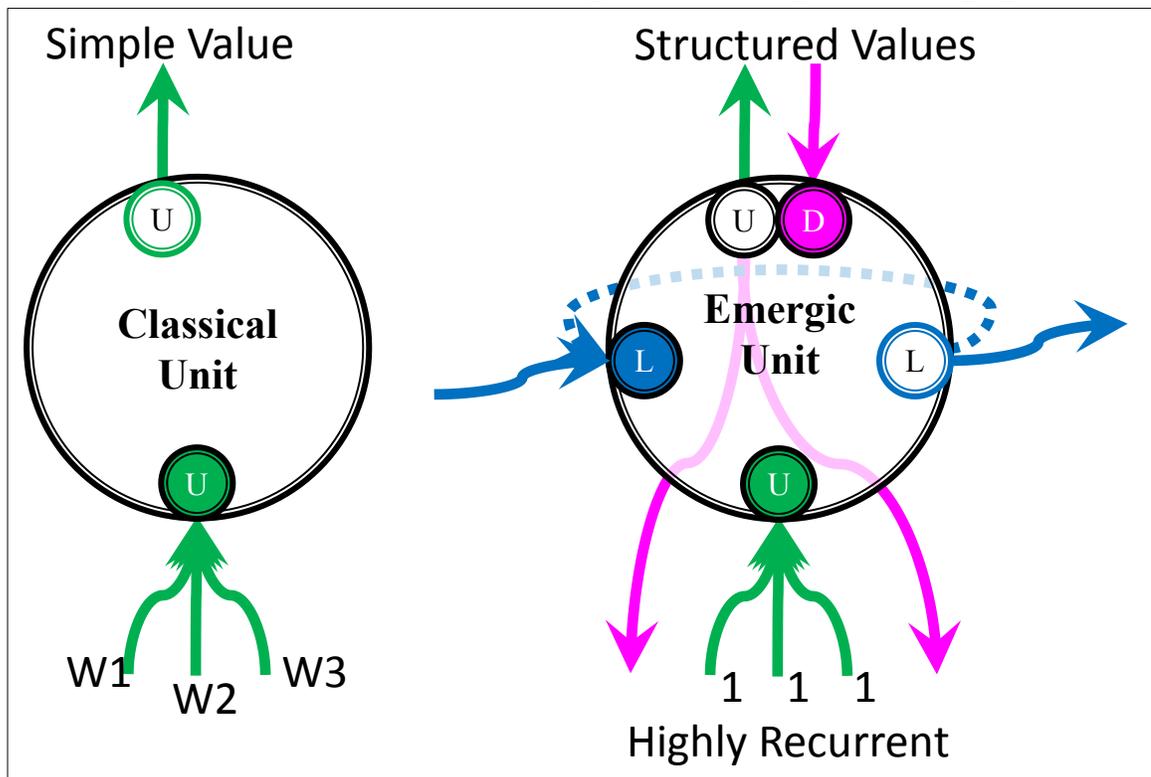
Illustration 2: Single feedforward (Up) flow in a classical unit vs. multiple interacting (Up, Down & Lateral) flows of an Emergic unit

when compared to the traditional storehouse metaphor. Also, this flow is engineered to be invariant to eye motion, so visual content is stable and does not pile-up to cause smear. A flow is a species of dynamic systems, but we demonstrate visually how it simplifies the understanding of dynamics for visual stability (e.g., see Illustration 6 on page 8).

Of direct significance to neurobiology, is that ECM suggests a possible mechanism for the phenomenon of receptive field remapping (Merriam et al., 2007), and a functional reason and context for its occurrence.

1.2 Emergic cognitive model

The Emergic Cognitive Model (ECM), when considered in detail, is novel in many respects. This makes justification extremely challenging, especially as its induction is based on *holistic* analysis. We attempt to avoid this problem as
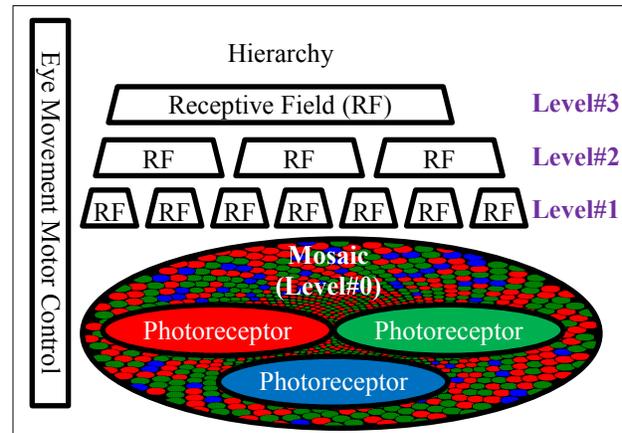
Illustration 3: ECM Structure

follows. Firstly, in this section we summarize ECM at a high enough level of detail such that the model does not appear that unusual. Secondly, the next four chapters and the next four appendices validate the model in action. These eight sections can be thought of as forming the *holistic* analysis for justifying our novelty. Finally, the novel details of the Emergic Cognitive Model (ECM) are found in Appendix G: (p323) without any rationalization, although there are considerations made to biological plausibility.

Here, we summarize ECM starting with a structural overview as depicted in Illustration 3. ECM consists of a sensory layer of cone photoreceptors, followed by a three level hierarchy of receptive fields (RFs) that are recurrently connected to allow for bottom-up, top-down, and lateral flows of information and interactivity. Flowing up, sensory information undergoes statistical spatiotemporal summation – emergic values can be

highly structured and in this case represent descriptive statistics tagged with spatial location. Flowing down allows for some interpolative (but not extrapolative) filling-in (as discussed on page 57 after tests results). Flowing laterally provides for a generative kind of memory and can impact filling-in to some extent. An eye is situated in front of a computer screen and is expected to engage with the image, i.e. saccade across the image. Eye motion is handled by a motor planning system that distributes motor instructions to all RFs in advance so that compensation *within the flow* can occur in time – we do not have spatiotopic neurons, only spatiotopic flows.

ECM realizes only two cognitive functions that interact. The *handling missing data* function is simplest and operates within the context of each receptive field (RF) as shown in Illustration 4. If an RF is missing information flowing up (⬆), then use lateral information if available (➡), otherwise use information flowing down (⬇). Then distribute that single value (if any) upwards, downwards and laterally (via Out).
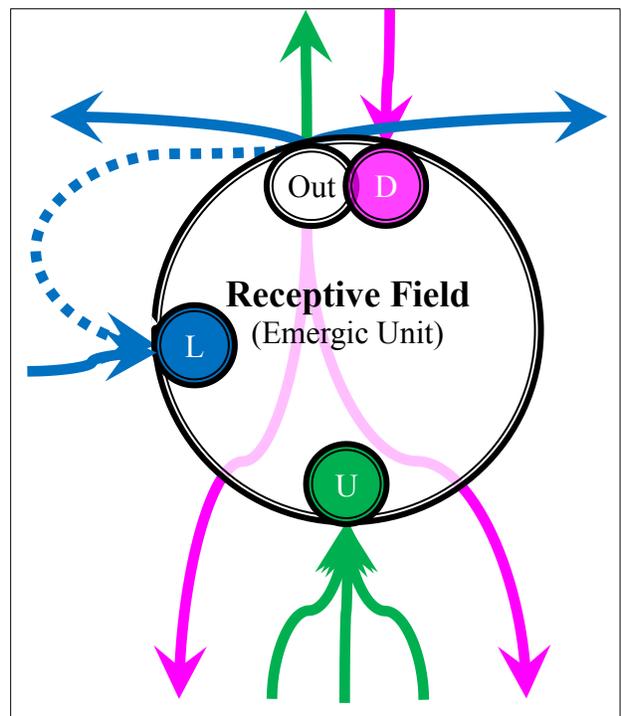
Illustration 4: Context for handling missing data function

The *maintaining information coherence* function is more complicated. In summary, it ensures that when *old* information about a spatial part of an image is sent upwards, and later reflected back down (or laterally), then when the reflection meets *new* information from that spatial source, they refer to the same spatial location *even* if the eye has moved in the meantime. For example, in Illustration 5, all the information about cyan will cohere in the next tick at the 2nd Receptive Field Layer if the downward flow from the 3rd layer is shifted by 1 unit leftwards and the lateral flow from the 2nd layer is shifted by 2 units leftwards. The upwards flow is never shifted. In practice, this means that if the eye saccades 3′ to the left, then information must be *relatively* shifted 3′ to the right to compensate, taking previous shifts into account. Of course, the shifting of information is complicated by the fact that shifting amounts cannot respect integral cone units, with RFs being heterogeneously distributed and of varying size. Information is statistically shifted in a manner akin to broadcasting a value locally as in a local area network with each RF only sampling the information within its new spatial purview – information values such as L, M and S activations are tagged with their spatial position.
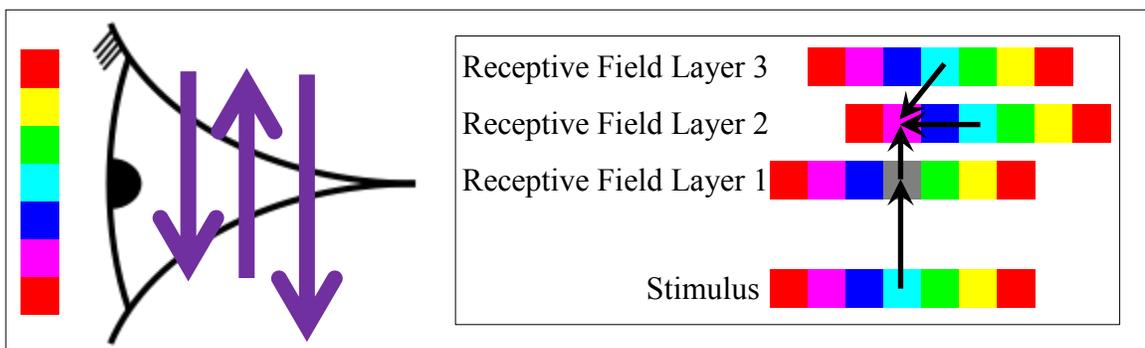
Illustration 5: Shifting down and lateral flows so cyan can fill-in coherently

1.3   Flowcentric vs. neurocentric analysis

In traditional *neurocentric* analysis, visual content and memory is attributed to individual neurons. Thus both retinotopic and spatiotopic reference frames (i.e., allocentric & egocentric) are looked for at the locus of neurons. On the other hand,
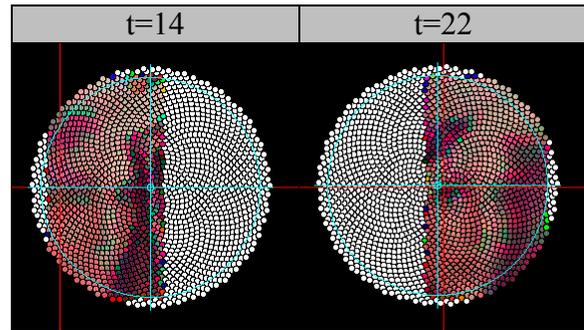
Illustration 6: Flowcentric stability

flowcentric analysis is a species of dynamic systems which attributes visual content and memory to the flow of information. In ECM, visual content is spatiotopic and stable with respect to the flow, but the flow can shift over the underlying neurons. This is shown within Illustration 6.

2   Chapter: An emergic model of filling-in for colour homogeneity

The *isomorphic* account of the filling-in phenomenon requires an *explicit* filling-in neural mechanism that, for example, takes colour information from border contrast or surround and diffuses it across a homogeneous surface (Boring, 1936; Pessoa, Thompson, & Noë, 1998). Such an explicit diffusion mechanism stops when it reaches other borders delimiting the end of the surface being filled-in. Computational accounts of filling-in naturally have analogous mechanisms to manage borders, surfaces and diffusion. Such accounts work well with artificial black & white stimuli.

However, in actuality, there are no crisp edges or well defined surfaces. At a low-level of neurobiological detail, the visual system is heterogeneous throughout – it does not have a regular pixilation as found in a digital camera; natural images have numerous colour gradients and rounded corners (Mather & Smith, 2002); and sharp borders become blurred under motion or due to differences in depth of focus (Risner & Gawne, 2009). Thus existing computational models may use a particular set of parameters to handle one specific case, but they may fail in other cases when using the same parameter settings – they are not robust (see Hu et al., 2011).

The Emergic Cognitive Model (ECM) has a single level of realization, so it must unite low and high levels of details that would ordinarily be investigated under different domains of inquiry. It is a parameter-free model that works in all the aforementioned cases specifically because it does not have surface, edge or diffusive mechanisms. These behaviours emerge when required, but more importantly, do not emerge when not

required. Thus the term *filling-in* in our account is descriptive, emergent and not explanatory – an isomorphic mechanism would always fill-in.

In this chapter, we focus on the low-level details of heterogeneities. We show how ECM can compensate for *neurobiological* irregularities by filling-in the sensory absences in a spatiotemporal manner, but will not fill-over *stimuli* irregularities, especially under motion.

2.1 Introduction

Many have wondered how the visual system can *perceptually* homogenize information from a heterogeneous *sensory* system (Hartridge, 1947;
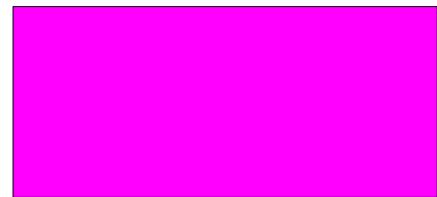
Illustration 7: Spatially homogeneous chromatic patch

Nakayama, He, & Shimojo, 1995). Compensating for imperfections is a classic visual problem (O'Regan & Noë, 2001). For example, why does a homogeneous patch of colour as in Illustration 7 look homogeneous when it does not "appear" so to the retina?

2.1.1 Reasons for Lack of Computational Models

Despite the simplicity of the question, there are few computational models available. Indeed, the "neural mechanisms of surface perception are surprisingly poorly understood" (De Weerd, 2011). There are many reasons for the lack of models.

Firstly, there are many different kinds of heterogeneities Nakayama et al. (1995), and they are distributed throughout the visual system, so corrections in one place may not be perceptually sufficient.

Secondly, while the stimulus parameter we are experimentally manipulating is surface colour, it is unclear whether other functional parameters are involved cognitively. For example, some models of surface perception require a prior stage of filling-in to occur from edges, borders, or contours (Grossberg, 2003). This implies that edge irregularities and false contours need to be spatially corrected beforehand. However, such low-level irregularities are glossed over, often because the genesis of the models comes from image processing where there were no sensory irregularities. Conversely, other filling-in mechanisms implicate a set of surfaces at many layers of stratification relating to border ownership and transparency (Nakayama et al., 2009). Each hypothesized system investigates their own part of the equation glossing over the sensory irregularities covered elsewhere. No model accounts for the overall equation with the simple question and stimuli we happen to use – why an actual homogeneous stimulus without any apparent borders, edges, layers, etc., looks homogeneous despite numerous sensory irregularities that could produce numerous border, edge or layer artefacts.

Thirdly, the homogenization function sits between *sensation* and *perception* implying that conscious awareness is somehow involved, and this certainly has no agreed upon bridge locus (Teller, 1984) to single neurons (Barlow, 1972; Thorpe, 2009) and may be impenetrable (Pylyshyn, 1999a). This issue of consciousness also relates to the second. Ideally surface colour perception would only need to be homogenized in a hypothetical last stage of visual processing before reaching conscious awareness, but again, no such last stage is known, and the concept of a last stage implies a Cartesian theatre viewed by a conscious homunculus (Dennett, 1991). If consciousness is somehow *distributed*

within recurrent brain dynamics (Taylor, 2007; Thompson & Varela, 2001), perhaps irregularities *do* indeed need to be corrected everywhere. On the other hand, patients with lesions in V1 can still report visual qualia indicating that the primary visual cortex (V1) and feedback to it are unnecessary (Ffytche & Zeki, 2011). Without a handle on consciousness, it is simply unclear how extensive and where the homogenization function must be.

A consequence of the problems related to consciousness is whether any kind of *isomorphic* filling-in actually occurs at a neural level. In the case of the blind spot, there are receptive fields over the blind spot (Awater, Kerlin, Evans, & Tong, 2005) so one can assume there must be some neural mechanism of filling-in. For brightness and color filling-in, there is no fMRI evidence in early visual cortex (Cornelissen, Wade, Vladusich, Dougherty, & Wandell, 2006). The lack of fMRI evidence for this fast version of filling-in could be because it is happening throughout a spatial layer all of the time, and perhaps because fMRI does not have the temporal resolution required. For slow types of filling-in there is some evidence (Hsieh & Tse, 2010), but this could just as equally be attributed to top-down attentional feedback.

Finally, most computational models that include biologically realistic colour cone

mosaics in their mechanism or take colour pathways into account, nevertheless have

trouble when showing colour, and may still use black and white stimuli or display their

visual processing results in grey-scale despite the fact that other figures and illustrations

in their paper are in colour. For example, in Illustration 8, Momiji et al. (2006) try to

show the negative filled-in colour of afterimages using a grey-scale rendition. The

obvious problem with this depiction is that without colour, it is not clear that the

afterimages are the correct complementary hue (which differs from the opponent hue)

(Pridmore, 2008a). More problematic is that it can clearly be seen that the patches of red

or green cones give the image a mottled surface, whereas the actual stimuli are

homogeneous. While the mottled surface is a sensory affair, the afterimage, and colour

homogeneity are both perceptual affairs.

Models with colour cones are historically about darkness/lightness perception and not

about colour. For example, the Feature Contour System (FCS) within Grossberg models

uses regions of continuous brightness for the perception of surfaces (Kokkinos, Deriche,

Maragos, & Faugeras, 2004). Thus a green surface beside a red surface of the same
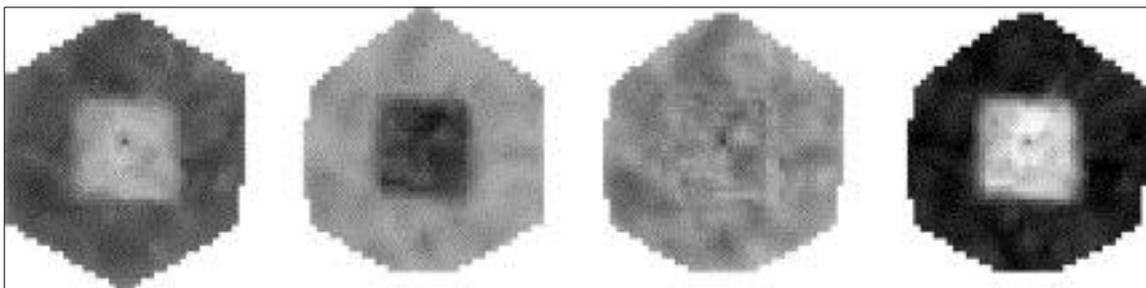
Illustration 8: Artifacts from grey-scale
From (Momiji et al., 2006): Fig. 19. Modelled foveal representation of **chromatic** images at 500 ms. The
15 × 15 square and its surround are composed of iso-luminance colours of λL = 565 nm and λM = 530 nm,
respectively. A single **grey-scale** is used.

brightness would be perceived as one surface; and it is the amount of brightness that fills in, not the colour.

2.1.2   Spatial irregularities to be homogenized

In this section we summarize the spatial irregularities that are obvious within the retinal photoreceptor mosaic and which we intend to homogenize with ECM. These are shown in Illustration 9. Similar irregularities involving numerous stimulus parameters can be found in receptive fields throughout the visual processing hierarchy. Some of these too will be homogenized. However, the many other irregularities are not considered at this time, because there are no counterparts within this initial simplified version of ECM.
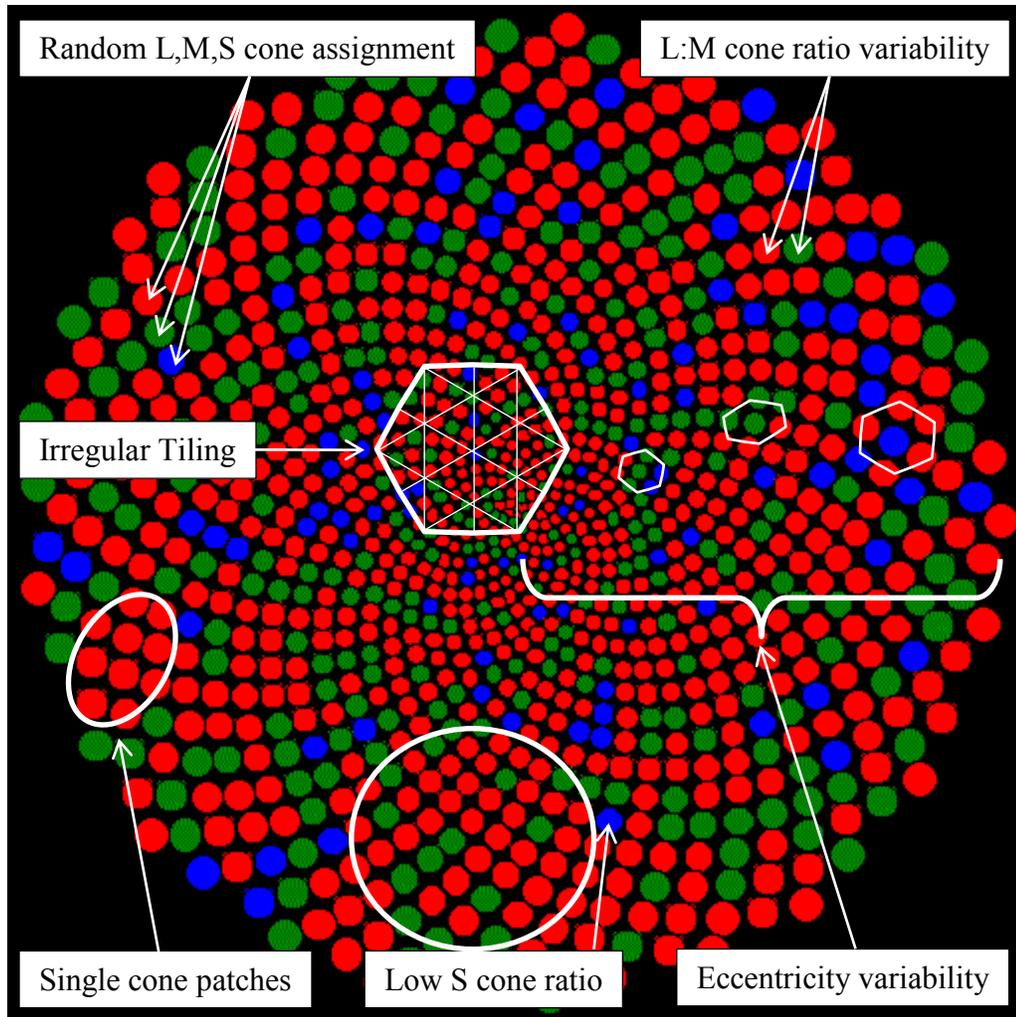
Illustration 9: Schematic of typical retinal mosaic irregularities

The basic irregularities are as follows. 1) Cone diameters increase with eccentricity

(visual angle away from center) so their density goes down (Curcio, Sloan, Kalina, &

Hendrickson, 1990). 2) Cone centers are not perfectly placed. Together, these lead to 3)

irregular cone tiling that could include some cones within only 5 neighbours or 7 rather

than the hexagonal 6 (Wojtas, Wu, Ahnelt, Bones, & Millane, 2008). There are, of

course, local regions of hexagonal tiling broken up by blue (S) cones that are

considerably larger than the others (not illustrated). The assignment of L,M,S sensitivity

is 4) random (Roorda & Williams, 1999). Along with 5) large individual variation in L:M

cone ratios while maintaining normal visual perception (Hofer, Carroll, Neitz, Neitz, &

Williams, 2005), this leads to 6) patches of cones with only one cone type where the eye is colour blind (Roorda, Metha, Lennie, & Williams, 2001) and where false contours could manifest. In Illustration 9, there were more L (red) cones, so most of the patches were of red. There could still be patches of M (green) only cones, but fewer in number. Patches could be greater than 5′ (Roorda & Williams, 1999). 7) S cones are always less numerous, leading to 8) solitary S cones and 9) large patches devoid of S cones (local tritanopia). It should come as no surprise that the filling-in of blue cones will be found to be problematic.

The irregularity that is apparently due to the optical blind spot (p161), and the fact that the foveoal center is devoid of S cones is handled in later chapters (p231). While we are only concerned with colour vision, we should point out that there are no rods in the foveola, and the numbers of rods also increases with eccentricity, so this too adds to irregular tiling. Temporal irregularities such as eye blinks are also handled elsewhere (p263). There are numerous other spatial irregularities throughout the visual processing hierarchy that are not covered. These include per neural class tuning heterogeneities (Schummers, Mariño, & Sur, 2004), great diversity in interneuron behavior (Isaacson & Scanziani, 2011) as well as with gap junctions (Arai, Tanaka, & Tachibana, 2010; Fukuda, 2007), dendritic structure (Meijering, 2010) and connectivity in general (Borra, Ichinohe, Sato, Tanifuji, & Rockland, 2010; Ko et al., 2011; Voges, Schüz, Aertsen, & Rotter, 2010), cortical maps with spatial irregularities that manage multiple stimulus parameters in varying amounts (Soo, Schwartz, Sadeghi, & Berry II, 2011), etc.

2.2 Review of literature and alternative models

We are not aware of any *specific* computational models for spatial colour homogeneity, so here we review the next best ones. Some of the models are conceptual rather than computational; some are based on image processing needs and may not be biologically plausible; while other explanations simply redescribe the main phenomenon in terms of another.

2.2.1 Problems with standard high-level account of filling-in

The standard *higher-level* account of filling-in is based on the fact that most information from the retina to the rest of the brain is not of surface properties but of luminance and colour contrasts. For example, Friedman et al. (2003) found that a homogeneous 4° stimuli will affect the high-level averaged V1 and V2 population response as shown in Illustration 10. This does not include the low-level heterogeneities that are the focus of this chapter and are obscured by

Illustration 10: Spatial population response profiles to homogeneous stimulus of 4°
From (Friedman et al., 2003) Fig 12: Average normalized response profiles of V1 and V2 neurons. Top, orientation-selective cells; middle, non-oriented cells; bottom, combined activity of oriented (filled) and non-oriented (dotted) cells.

averaging. While this could simply be a matter of encoding, there does exist an actual perceived phenomenon of filling-in that takes a noticeable amount of time to complete (Huang & Paradiso, 2008; Murray, Foxe, Javitt, & Foxe, 2004; Paradiso & Nakayama, 1991), although some can be very fast (Robinson & De Sa, 2008). Therefore, it is

hypothesized that the perception of surface colours is due to a high-level filling-in

process that causes edge contrast information to flow across a uniformly smooth area

(Arrington, 1994) till it encounters a real (or illusory) contour and then stops (Feitosa-

Santana, D'Antona, & Shevell, 2011). The problem for our lower-level of analysis is that

due to heterogeneities, there are patches of single cone colours, and every contrast

detector has a different ratio of colour cones in their surround. At the low-level, even a

homogeneous stimulus appears to be full of edge information – there is no "uniformly

smooth surface" as that is precisely what we are trying to explain. These irregularities

that prevent a smooth surface as appearing uniform would have to be effectively

homogenized before the standard high-level account of filling-in could apply – otherwise

the high-level filling-in process would stop immediately as it encountered a discontinuity.

The other problem with the high-level account is that it is related to Troxler's fading

(Troxler, 1804). The edge detecting neurons must first adapt before filling-in from the

surround would occur. However, the low-level colour homogenization happens

instantaneously without waiting for anything to adapt. There is also a conceptual gap in

the explanation and the phenomenon. If the edge neurons adapt, then they may not have

the information required to fill-in, although the surround *is* available for this purpose. In

any case, it is not explained at a conceptual level why the surround would fill-in rather

than the interior fill-out as can happen (Breitmeyer & Jacob, 2012; Shimojo, Wu, &

Kanai, 2003). Indeed, it is the very need for hard-coding of white being the surround

while black being the interior to be filled-in that makes Arrington's (1994) lightness

filling-in model incompatible with the opposite darkness filling-in (Rudd & Arrington, 2001).

(Hamburger, Prior, Sarris, & Spillmann, 2006) show that the filling-in/out direction is related to what is considered as figure and what as ground, implicating top-down object vision and attentional factors. Perhaps the hard-coding of black or white figures is unintentionally a proxy for these top-down influences. Nevertheless, in natural scenes, it is difficult to argue that attention is differentially allocated to the multitude of irregularities at a glance.

Arrington (1994) uses the Boundary Contour System/Feature Contour System (BCS/FCS) of Grossberg et al. (Cohen & Grossberg, 1984; Grossberg & Todorović, 1988; Grossberg, 1983) for a computational model of higher-level filling-in. Despite the mention of colour contrast information, the stimuli are black & white, and the 1994 paper is about brightness filling-in. While Arrington assumes that darkness filling-in would be similar, he did not model this, nor could ask questions about possible interactions related to top-down mediated figure/ground segregation as we have seen. A separate darkness filling-in model had to be constructed (Rudd & Arrington, 2001) which has grey-scale colour constancy abilities borrowed from Retinex (Land & McCann, 1971; Land, 1986) – filling-in signals are not stopped at edges but are partially gated. Extrapolating from this disjointed modelling approach (one model for lightness induction, another for darkness induction), is that with respect to colour, every hue pair combination might need a different model or process. One reason to remain in the grey scale world is that the

achromatic pathway only needs to be used. It equals L+M (Red+Green) and is thus invariant to many of the heterogeneities we are concerned about – whether a cone is red or green makes no difference. The S (blue) cones do not contributed to the achromatic pathway and are so few as to not make patches.

2.2.2 Filling-in across edges

Ryota Kanai et al. (2006) have shown that unlike the standard filling-in accounts where filling-in is stopped by real or illusory contours (Feitosa-Santana et al., 2011), there are cases where filling-in can cross edges, i.e., the gridlines of Illustration 11, based on what must be a bottom-up and top-down "integrative
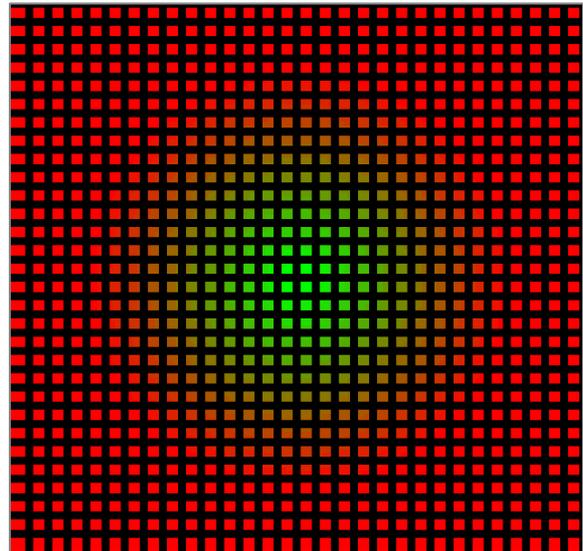
Illustration 11: Filling-in across edges.
After (Kanai et al., 2006)

computational scheme" related to the grouping and segregating of elements and surfaces. This filling-in occurs without the gridlines fading-away – it is the blurred border of the green disc that takes time to adapt prior to the filling-in (or out).

Perhaps the same mechanism crosses the heterogeneities of the retina and can fill-in the patches? One problem with this solution is that the gridlines or edges due to heterogeneity itself are not filled-in while we simply do not perceive any heterogeneity when looking at a homogeneous patch. The second problem is that filling-in must await edge adaptation whereas homogeneity is immediate. The third problem with this account

is that it explains a part of the main phenomena with another – there is no computational model of this crossing the gap effect.

2.2.3 Induced filling-in

The filling-in for homogeneity happens immediately and cannot await edge adaptation. Simons et al. (2006) found that they could instantaneously induce a fading of complex stimuli across the periphery and fovea via the presentation of visual transients. The image was filled-in (homogenized) with the *average* scene hue and luminance rather than being replaced with surround statistics as with standard filling-in accounts. This implies that filling-in was not by propagating edge contrast information as would occur in standard accounts, but from entire scene statistics. The interesting aspects of this phenomenon, beyond implying a top-down mechanisms to trigger the fading/filling-in process, is that the filling-in was not stopped by scene borders as would occur in the standard filling-in accounts – although it should be pointed out that the scene was blurred to begin with so as not to have crisp borders. Due to cone heterogeneities, there are numerous patch borders, but these do not stop the filling-in of a homogeneous colour. Thus this induced filling-in phenomenon could apply to our homogenization phenomenon, but it is simply a relation to or redescription (Gigerenzer & Brighton, 2009) in terms of another phenomenon lacking a computational model, and it is not likely that attention is involved in the same manner.

2.2.4 Averaging effect of Diffusion

The averaging mechanism of induced filling-in is of interest as it could average out the heterogeneities. In Grossberg & Todorović (1988), the filling-in is a diffusion process which does tend to average out the brightness precept as would be useful for brightness

constancy, however it is stopped at borders, and the averaging value comes from the edge

or immediate surround, but not from the entire scene. Diffusion doesn't eradicate the

heterogeneities as would be required in our homogenization phenomenon. Instead, it

tends to blur the image which is a common artefact of filling-in models (Keil, Cristóbal,

Hansen, & Neumann, 2005).

D. Hu, Zhou, & Wang, (2011) demonstrate how Grossberg's FACADE model can be

used in an engineering application to

remove the heterogeneities associated

with noise. While Illustration 12 shows

that these heterogeneities are

homogenized over, the problem remains
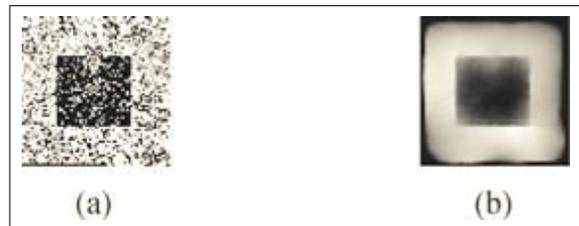
Illustration 12: Noise Removal using FACADE.
From (Hu et al., 2011) Fig. 8: (a) Test image with noise; (b) FACADE processed with typical model parameters.

that human subjects have the acuity to see this salient noise or texture. In this particular

example, the edges of the square stimuli are crisp as compared to traditional filtering

accounts which produce artefacts. That is due to the boundary completion system (BCS)

above and beyond filtering.

2.2.5   Regularization Theory

An extreme example of engineered vs.

cognitive applications is seen in

Illustration 13 where half the image is

restored via regularization (Poggio, Torre,

& Koch, 1985). This includes the removal

Illustration 13: Extreme homogenization
From (Satoh & Usui, 2009) Fig. 7: a) Stimulus with half visual information missing; b) Restored response.

of the other half that is, nevertheless, perceptually salient. While this regularization

scheme may be biologically informed by V1 and V2 physiological evidence, it is not

phenomenological plausible at such an extreme scale. In particular, any regularization

scheme (including the brain) requires *a priori* knowledge which results in differential

artefacts often made apparent by illusions. While each algorithm may excel in some

dimensions, such as the filling-in as illustrated, there may be numerous other artefacts

(such as impact to acuity), that are not shown in the careful selection of processed

images. Without a unified approach across a range of stimuli, individual algorithms such

as (Satoh & Usui, 2009) used above are impossible to evaluate on a cognitive basis. They

are examples of biologically inspired but no longer plausible models.

Note that almost all engineered systems use a normalized and regular RGB mosaic, and it

is not clear that any specific type of added noise emulates the heterogeneities involved in

this paper. Some of these heterogeneities have been modelled, as in (Deering, 2005;

Momiji et al., 2006; Wachtler, Doi, Lee, & Sejnowski, 2007), but for other purposes – not

for colour homogenization. Heterogeneities can be functionally important and worthy of

engineering. For example, one consequence of the quasi-random spatial sampling is

acuity better than that implied by the Nyquist rate as aliasing and moiré effects are

scattered in broadband noise (Yellott, 1983). This is simply not available on regular

sampling.

2.2.6 Cone type reconstruction

Although receptive field surrounds are a heterogeneous mixture of L and M cones bound

by a local L:M ratio, it may be possible under certain assumptions to reconstruct the

trichromatic underpinnings (hue and position) via an unsupervised learning algorithm

based on natural scene statistics to allow for selective rewiring during development (Doi, Inui, Lee, Wachtler, & Sejnowski, 2003; Wachtler et al., 2007). Perhaps such a reconstruction can be used for homogenization? Unfortunately this account was not proposed for such purposes, and suffers from many recognized problems (Wachtler et al., 2007) making regularization impossible. Firstly, in the cases where the assumptions are valid, the account at best recapitulates the original heterogeneities so offers no mechanism for normalization. Secondly, there is a loss of acuity which increases image blur. Thirdly, the current algorithms for reconstruction based on independent component analysis (ICA) are not biologically plausible due to non-local learning rules (Wachtler et al., 2007), but see Lin, Grier, & Cowan (1997). Mechanisms with local learning rules have been proposed. A Bayesian model provides a good first order account for stimuli consisting of small spots (Brainard, Williams, & Hofer, 2008). It is based on optimality considerations rather than known biological mechanisms.

2.2.7 Adaptation and selective wiring hypothesis

The visual system is filled with adaptation mechanisms that can function as gain control useful for normalization (Solomon & Lennie, 2005). For example, heterogeneities involved with entoptic images, e.g., retinal blood vessels, can be compensated for extremely quickly (Coppola & Purves, 1996). Why cannot our cone mosaic heterogeneities be similarly compensated for? In one respect they can. When adults are fitted with coloured

Illustration 14: Mixed-wiring in cone opponency.
From (Demb & Brainard, 2010) Box 1 showing random cone opponent wiring in fovea and periphery.

filters, a shift in perception occurs which can last for two weeks after removal (Neitz,

Carroll, Yamauchi, Neitz, & Williams, 2002). This effectively normalizes the average neural or opponent response mechanism to global scene statistics, but it does not imply any filling-in of heterogeneities across space. For example, each receptive field may be normalized such that the long-term global average scene statistics represents white (Webster, 2009), but if a green stimuli is shown, then within red cone patches, there will not be any filled-in signals representing the missing green cones.

Adaptive mechanisms work by affecting the gain of some neural signal. The problem is that after different combinations of L, M and S signals are multiplexed into a receptive field, then a single further gain could not tease the original contributions apart whereas reconstruction would. Certainly, the individual L, M or S cones can be individually adapted prior to multiplexing. This could be used to have every high-level RF have equal amounts of each cone type, much like an RGB pixel on a computer monitor. But this does not occur. Adaptation could ensure center-surround antagonism is pure, e.g., assign a gain of zero to a green photoreceptor in a surround of mostly red, e.g., the top left RF in Illustration 14, but this would be equivalent to the selective wiring hypothesis which ensures purity by wiring alone (Demb & Brainard, 2010). Selective wiring does not appear to occur in the fovea, although in the periphery, there may be some selectivity in RF centers (Hansen, Pracejus, & Gegenfurtner, 2009). Without some selectivity in the periphery, the center and surround would have similar amounts of L & M which minimizes opponency and makes red/green discrimination difficult as is found empirically (Ferree & Rand, 1919; Hansen et al., 2009). The random or mixed wiring hypothesis has the most support, especially in the fovea (G. D. Field et al., 2010; W.-C.

A. Lee & Reid, 2011). However the debate continues (B. B. Lee, Shapley, Hawken, & Sun, 2012) and it may well be that feedback from a higher level homogenization function provides the physiological evidence for some selectivity when it cannot be found anatomically (Crook, Manookin, Packer, & Dacey, 2011). In this case, such evidence would represent a methodological problem, as feedback was not taken into account within a feedforward conception of wiring.

Adaptation, therefore, does not help with the filling-in of heterogeneous LMS mixtures across space. For the fovea, the center of a center-surround RF is always pure being a single cone, and the surround is on average similar to the local L:M ratio, thus simulations indicate selective wiring is not required and does not match physiology (Momiji, Hankins, Bharath, & Kennard, 2007). This does not help the heterogeneities around patches – the non-average cases. Moreover, despite similar opponencies on average, this cannot correct for the heterogeneous *position* of cones. For example, no amount of wiring or adaptation can add the missing S (blue) cones back into the blue-free region of the foveola. Similarly, any large patch of single cone colour, even if rewired to further off cones, would have the additionally wired cone types express a wrong location which would affect acuity. More importantly, while adaptation affects gain, it cannot make cones sensitive to a new location which is ultimately what would be required for any spatial regularization.

Very few cortical neurons are involved in chromatic tuning as this would typically affect spatial tuning (Solomon & Lennie, 2005). The acute density of the fovea, the lack of blue

cones, the historical lack of M cones (Surridge, Osorio, & Mundy, 2003), suggest that
spatial information was of highest priority and selectivity for colour less so. Adaptation
may normalize average colour response, but cannot regularize space.

2.2.8   Pan-field colour

When a colour image is modified to
contain achromatic regions, as in
Illustration 15, participants often judged
them to be completely full of colour. The
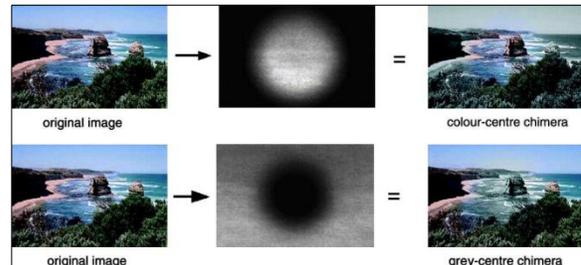
Illustration 15: Pan-field colour
From (Balas & Sinha, 2007) Fig. 1

interesting aspect of this pan-field colour phenomenon is that, it is based on a fast
completion mechanism, and it is a mechanism that does not stop at edge contrasts. This
same mechanism could also be used to fill-in the missing colours of heterogeneous cone
patches missing one or two cone types. While the standard filling-in account is based on
diffusion from an edge inwards, this account could be a top-down coarse-to-fine process
(Hegdé, 2008) based on global image statistics – the gist, or visual short term memory
(Velisavljević & Elder, 2008), or of recognition memory (Potter, 1976) associating
colours to objects (Oliva & Schyns, 2000). Alternatively, it may be that the detecting of
achromatic regions is made difficult because the presence of global scene statistics make
us blind to certain changes (Sampanes, Tseng, & Bridgeman, 2008), in which case
attentional aspects may confound the experimental design.

One of the great debates for any filling-in phenomenon is whether it is real or an illusion
– are we *filling-in* beforehand or *finding-out* afterwards (Dennett & Kinsbourne, 1992;
Dennett, 1991, 1998). For example, the pan-field filling-in may be an example of an

illusion – the achromatic region is not noticed, but when asked, it is filled from memory of similar objects or nearby gist, but not based on any neural process completing the chromatic field from actual sensory data. If illusive, it is no wonder that this colour completion can appear to "cross" borders – there is no isomorphic function during sensory processing.

On a similar note, in the preceding section on random wiring we noted that the periphery becomes increasingly achromatic as both center and surround are large enough to have a common L:M ratio within and without (Ferree & Rand, 1919; Hansen et al., 2009). Dennett (2001) notes that perceivers are surprised when this is demonstrated to them – they were not aware of their peripheral loss of colour vision in their entire life (nor their loss of resolution).

If pan-field filling-in is a real mechanism, i.e., with an isomorphic brain mechanism, then because no computational model has yet been developed to explain this phenomenon, it does not help to redescribe our homogenization phenomenon in terms of it.

2.2.9   Colour constancy

The phenomenon of colour constancy is normally associated with perceiving the same colour under different *illumination* – a problem of recovering a unique surface spectral *reflectance*. For example, in Illustration 16, the blue tiles on the left (under yellow light) and the yellow tiles on the right
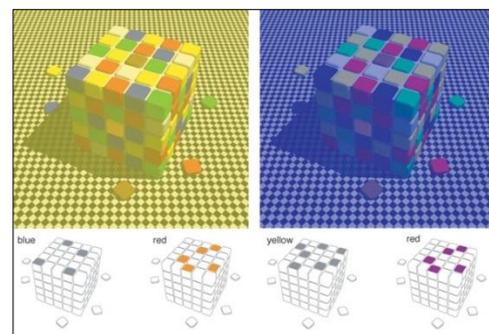
Illustration 16: Colour constancy under illumination variation.
From (Lotto, Purves, & Nundy, 2002)

(under blue light) are identical shades of grey. Thus while blue can be *perceived* as constant under various illuminants, it may be that grey is what is *sensed* in a particular case. Could this form of filling-in explain homogenization?

While there are numerous models for this form of colour constancy, they all effectively operate over a larger region of stimulus surface. They will thus preserve the smaller scale irregularities we wish to homogenize. Constancy typically retains contrast ratios (Pridmore, 2008b), especially across borders. For example, a red patch of cones next to a green patch will still retain similar red/green contrast under various illuminations. Indeed, one can imagine that the tiles on the cube are such patches – and they have not been homogenized at all – simply equally altered. The problem, of course, is of adequate thresholds. Surfaces are found by the removal of gradual discontinuities that could indicate a shadow for example and then constancy adjustments are made (Kingdom, 2011). Any threshold large enough to include both red and green cone patches under one surface would include the entire image as one surface. Moreover, if one considers our heterogeneities as a form of texture or inconsistent structure, then this complicates current constancy analysis as a unique surface spectral reflectance is not well defined (Foster, 2011).

However, colour constancy is not a unitary phenomenon (Foster, 2003), and can involve invariance under a more pertinent kind of change. In Illustration 17, properties such as gloss are variable throughout, yet the blue cup is perceived as homogeneously blue at first glance (Brainard & Maloney, 2011). Perhaps it is this kind of filling-in that occurs to homogenize our cone heterogeneities? The problems here are several. Firstly, the account is similar to the pan-field colour account in that gist statistics are involved, and this may prove an illusion of constancy. Under detailed examination, the gloss artefacts are salient and can be perceived. This is
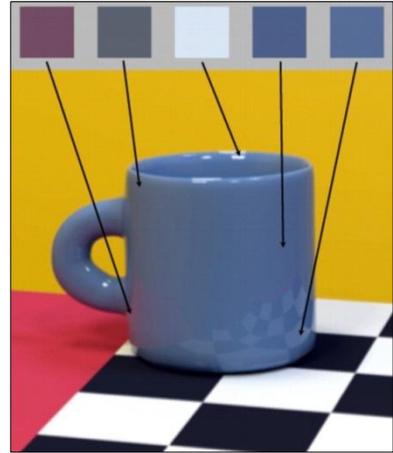
Illustration 17: Colour constancy under gloss variation.
From (Brainard & Maloney, 2011) Fig. 1: Individual locations on a coffee mug can reflect very different spectra. The uniform patches at the top of the figure show image values from the locations indicated by the arrows. The mug itself, however, is readily perceived as homogeneous blue.

unlike the homogeneous patch where cone heterogeneities cannot be perceived despite attempts to do so. Secondly, as far as removing gloss, there are no known computational ways to do so without *a priori* knowledge. Most computational filters would introduce blur, especially when required to remove large areas such as the white reflection on the top lip of the cup. The only reason a polarized lens can remove such artefacts (and darken the entire image) is because it has access to information about light beyond RGB image values. In general, there may be valid reasons why other regions of an image contain small areas of white, so removing such supposed artefacts could destroy the integrity of the image.

2.2.10   Sensorimotor account of "seeing" a patch of red

The standard phenomenological account of "seeing" a unique hue of red somewhere within a static picture relates the qualia of redness with the amount of L cone activation (in comparison to other cones), or the L:M opponency activation values (in comparison to other opponencies). One could call this the RGB account of colour. There is an alternative.

The sensorimotor account of "seeing" red relates the qualia of redness with active engagement or exploration of the scene (O'Regan & Noë, 2001). It is not the amount of redness at a point, but how the various activations within the visual system change as the eye moves across the point. In particular, the qualia of redness would include the dynamics of heterogeneous responses however distributed and there would be no need to homogenize the irregularities. As an example, D'Zmura & Lennie (1986) have shown that by coupling light adaptation to eye movements, hue can be estimated independent of object shape or position. Indeed, without eye movements, all of vision fades away (Ditchburn & Ginsborg, 1952; Troxler, 1804; Yarbus, 1967).

The sensorimotor account suffers from a significant conceptual problem – how does a dynamic property manifest itself in a static view of perception. Nevertheless such a conceptual shift is required (Broackes,
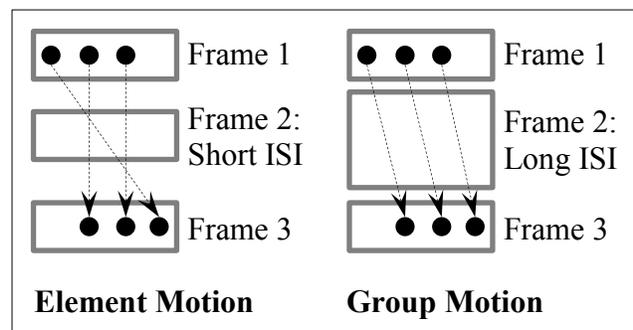
Illustration 18: Element or group motion percepts.
After (Aydın, Herzog, & Öğmen, 2011), Fig. 2.]

1992). Possibly a good example is to be found in the perception or qualia of apparent motion as happens in animation. Qualitative results can be measured using a Ternus-Pikler display (Pikler, 1917; Ternus, 1926) as shown in Illustration 18. When the interstimulus interval (ISI) is short between two frames, a single element is perceived to move; when the ISI is made longer, the entire group is perceived to have moved. These qualia of motion are assigned to the element or group, but are not available within any static snapshot.

The point, of course, is that redness need not be some unique hue constancy derived from a static image, but can be related to the dynamics associated with an engagement of red.

2.2.11 Attention

It may be that one cannot *attend* to sensory heterogeneities as they may fall below the nonlinear threshold for access consciousness (Del Cul, Baillet, & Dehaene, 2007). For example, in the case of hemispatial neglect "patients fail to attend to, report or represent information appearing in contralesional hemispace, despite intact sensory processing and visual acuity" (Marotta, McKeeff, & Behrmann, 2003). What is not attended to cannot be perceived, and this can be demonstrated in visual tasks as depicted in Illustration 19.
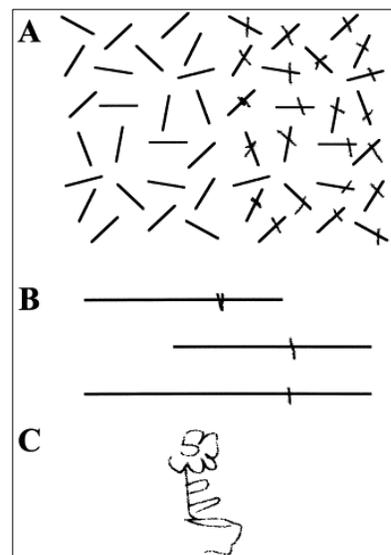
Illustration 19: Visual deficits with neglect.
From (Driver & Vuilleumier, 2001) Fig. 1: (A) Line cancellation task – mark all lines; (B) Line bisection task; (C) Draw from memory or copy from template.

For healthy participants, cognitive tunnelling and change blindness are a form of inattentional blindness that occurs regularly. Simons & Chabris (1999) demonstrated that a salient black gorilla thumping its chest and crossing the fovea is not perceived by 50% of the participants because they were engaged in a cognitively demanding task related to watching the activities of a white shirted team in a movie, one frame shown as Illustration 20.

Illustration 20: Inattentional blindness to gorilla
From (Simons & Chabris, 1999).

Thus, irregularities may not be perceived because of "our tendency to fail to attend to the heterogeneity and complexity of experience" (O'Regan & Noë, 2001). A previous example is found in Illustration 17, where the heterogeneous distribution of gloss is not noticed at first glance and the cup is perceived as homogeneously blue. In that particular case, attention can be focused on specific areas of gloss as to perceive those irregularities. While attending to irregularities in the world is easy, attending to irregularities of the visual system, such as the blind spot, is difficult – we are not meant to sense our visual system, merely to use it. The pan-field colour phenomenon may also be a result of attention.

While the attention account is related to the sensorimotor account – the qualia of redness includes associated attentional dynamics –it could hypothetically be a distinct homogenization mechanism, or an illusion of homogenization as nothing is actually filled-in nor perceptually sought after. Instead of isomorphic filling-in, we have

attentional modulations. Supposedly both spatial and temporal integration thresholds would need to be reached at that retinal location before an actual image edge or speck could be perceived to some degree of acuity. Unfortunately, this account is missing a computational model that involves attention, and so cannot be further compared against ECM.

2.2.12 Image Processing Demosaicing

In image processing parlance, our homogenization function is equivalent to demosaicing algorithms as used in handling various charged-coupled device (CCD) sensors within digital cameras, for example (Alley, Susstrunk, & Herault, 2005). Typically, the arrangement of pixels is regular, for example they are often *painted* according to the Bayer Color Filter Array which introduces aliasing artefacts (Hérault & Durette, 2007). These algorithms process one image at a time unlike the visual image processing next. They are noted here for completeness, but have not been used for irregular sensory mosaics.

2.2.13 Visual Image Processing Handling Missing Data

Defects within the retinal mosaic, such as colour patches missing one or two colour cones, scotomas missing all colour cones, etc. have counterparts in the video image processing realm where they may be called dirt, scratches, sparkle, speckle, blotches, noise, moiré patterns, tears, etc., with typical approaches addressing each issue separately. These corrective measures lead to a high-rate of false alarms, when motion across frames is poorly estimated. Bayesian approaches can unify motion, missing data, and noise jointly (Kokaram, 2004). The unified approach may be worse than dedicated algorithms for images containing specific problems, but the Bayesian compromise shows

overall improvement. The Bayesian algorithms within the image processing domain need not be biological plausible, but that does not mean that cognitive approaches cannot be inspired by engineering approaches. For example, if one frame of an image is missing information which is available from another part of a previous image due to object or camera motion, then the previous information could be used in a *cut & paste* manner (Van Roosmalen, Kokaram, & Biemond, 1999). The point is that many cognitive mechanisms are approached based on static image snapshots, while some engineering approaches are based on a world of motion, similar to the situation of actual cognitive agents.

2.3   Emergic model

In order to explain the phenomena of this chapter, we have taken the Emergic Cognitive Model (ECM) architecture as summarized on page 5, and made no changes.

In this chapter, we wish to characterize the relative contributions of the *handling missing data* and *maintaining information coherence* functions to the emergence of filling-in. In particular, we can translate our characterization exercise to the impact of lateral flows, top-down flows, and eye movements which can be precisely quantified. In this chapter, we minimize eye motion, but not jitter, so that lateral flows of information contribute minimally, while top-down flows contribute maximally.

2.4   Test 1: Colour patches viewed by L:M:S cone ratios of 7:3:1 under jitter

While all human retinas are heterogeneous as schematized in Illustration 9, those with L:M cone ratios exceeding 1:1 would contain the largest number of red only or green only cone patches. Moreover, the worse the ratio, the larger these patches tend to be. For

a human participant with a 3.79 L:M ratio, the patches can be larger than 5′, with colour

vision remaining normal (Roorda & Williams, 1999). Note that 20/20 vision corresponds

to acuity of 1′, and we assume the subject did not notice their 5′ patches – a homogeneous

patch will look homogeneous without any 5′ holes.

In this test, we wanted to characterize how the ECM architecture performs under a steady

(yet jittering) eye having all the aforementioned heterogeneities including a 7:3:1 ratio of

L:M:S cones.

2.4.1 Methods

The Emergic Simulation System (ESS) can engage both human participants and virtual

agents with identical stimuli (p365). Nevertheless, some methodological differences will

ensue. For example, while the L:M:S cone ratios for a virtual agent can be engineered, it

must be measured for a human participant. While the attentive and motor aspects of a

virtual agent can be strictly controlled, humans need additional instruction, and

preventing microsaccades in human participants may be more difficult. While the stimuli

are perfectly presented to the virtual agent on a virtual computer screen, a human

participant may need to be in placed in a dark room or one with diffuse light to prevent

shadows making the stimuli strictly nonhomogeneous[1]. Moreover, the virtual agent has a

minimal retina, so the stimuli would need to be made larger for human consumption.

Finally, while the internal structural or functional aspects of the agent are known and can

be displayed, human internal or external behaviour need to be measured with knowledge

---

[1] In practice, all static stimuli in this paper can be used as printed in colour under normal reading
conditions. Controlled conditions are not required for the level of detail we wish to uncover.

of neither system level structure nor function. Comparisons between the two must involve intuitions as to their relevance.

In this section, we indicate the methods applied to the virtual agent with sufficient information as to enable the specification of a corresponding experimental paradigm for human participants.

2.4.1.1 Subject analysis summary

A single cyclopean virtual agent was stochastically generated based on the development parameters detailed in the supplementary material. Its retina is shown in Illustration 21. Such parameters are intended to remain fixed across all tests; nevertheless, they may be changed

Illustration 21: Stochastic retina for homogeneity testing

to focus experimentation on one aspect at a time. For example, there is no blind spot in this agent. Such development changes are not considered to be of material consequence and simply represent subject variability. The cognitive model underlying the agent remains fixed and parameter free. The important parameters for this test are that the L:M:S cone ratio is 7:3:1 and that eye jitter is retained in this and all other tests.

2.4.1.2 Apparatus

All ESS tests produce stimuli on a computer screen with the agent a set distance away. We will not indicate this distance, but note the visual angles used. This allows translation to agents of different developmental parameters, and to human participants.

2.4.1.3   Stimuli

A single homogenous chromatic patch was
presented to the cyclopean agent at time zero. It
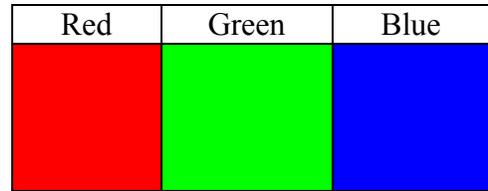
Red	Green	Blue

Illustration 22: Test 1 Stimuli

covered the entire visual field of view for the entirety of the test. Three test conditions

were tried – a red, green and blue patch. They are fully described in the supplementary

material (s3), and appear as in Illustration 22.

2.4.1.4   Procedure

At time=0, one of the red, green or blue stimuli was presented to the single eye of the

virtual agent. The agent was "instructed" not to move its eye, although this would not

prevent jitter which is under involuntary control. The artificial brain was monitored for

the first 16 time ticks.

2.4.2   Results and discussion

The visual frame-by-frame response results, with precise timing information, are detailed

in the supplementary material (s6, s12, s18). In this section, we summarize and discuss

various highlights. We also extract filling-in times.

The frame-by-frame account encompasses more than 36 pages because animations cannot

be printed. It is illustrated within tables that function as a multi-track sequence diagram

annotating the dynamic "orchestra of the mind". When presented as an animation

(Leibovitz, 2012a), all the results for all three conditions, along with Tests 2, 3 & 4, can

be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=31
```

The animations exhibit behaviour, e.g., diffusion and border processing, and can indicate movement that is hard to detect across image frames, while the frame-by-frame account allows precise quantitative comparison of processing over time. Here we highlight some of the data in a more useful form via print technology.

In all cases, each image frame within an animation or alone, represents the set of emergic values (p321) that have just been sent in the previous computation – technically emergic units, e.g., photoreceptors and RFs, do not retain state. The only exception is that when the ESS is not running, e.g., at tick 0, then the photoreceptors only, do not show what they last sent, but what the view currently looks like. This enables the operator to move the eye's fixation point and ascertain precisely what the eye is looking at.

2.4.2.1   Red filling-in for RF level 1

The 7:3:1 L:M:S cone ratio favours red with few green only patches. There were no blue only patches being just 10% of all cones. Thus red can be filled-in the quickest within 4 ticks (Illustration 23). While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.

t=2	t=3	t=4	t=5	t=6
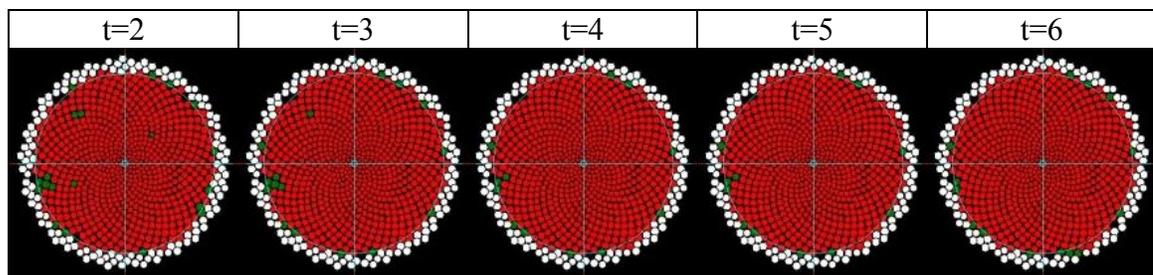
Illustration 23: Red filling-in for RF level 1

Note that we only consider the filling-in of the interior regions. We discuss the exterior later, shown as white RFs.

2.4.2.2   Red filling-in for RF levels 2 & 3

The higher level RFs do not require any filling-in as their spatial summation extent is sufficient, i.e., the size of the green only patches is too small.

2.4.2.3   Green filling-in for RF level 1

There are more and larger red only patches that need to be filled-in by green necessitating 12 ticks for completion of the interior as shown in Illustration 24.

t=2	t=3	t=4	t=5	t=6	t=7	t=8

t=9	t=10	t=11	t=12	t=13	t=14	
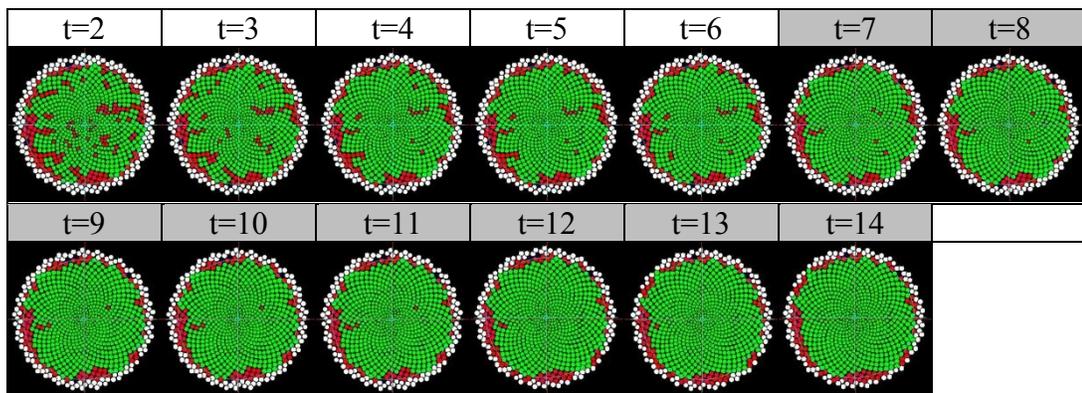
Illustration 24: Green filling-in for RF level 1

Note that at time=7 (as indicated by grey shading), the motor plans for eye jitter have finally caused the RF level to shift, and there is a marked increase in filling-in at this time. Both the shift and filling-in are more easily observed in the animation sequence. They are enlarged below in Illustration 25. The red cross-hairs are only visible at time=7

and represent shifted headcentric coordinates, while the cyan crosshairs and circles represent retinocentric coordinates and the boundaries of the two retinal photoreceptor regions – the foveola and fovea. In this case, we have minimized the

t=6	t=7
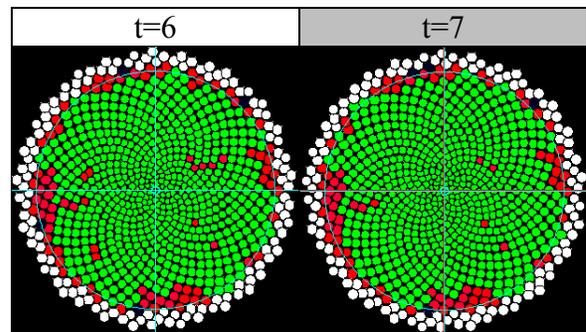
Illustration 25: Green filling-in for RF level 1 (expanded view after initial jitter)

foveola as it is ordinarily devoid of blue cones and we leave that to another section to

investigate (p231). What can be observed is that the eye has jittered both right and downwards by a small amount.

2.4.2.4   Green filling-in for RF level 2

The red-only patches are larger and so some level 2 RFs need to be filled-in with green as well. This takes 2 ticks as shown in Illustration 26. Level 2 RFs are larger than level 1 RFs from whom they receive bottom-up information, but after a one tick delay. Thus this level appears to start one tick later, at t=3. The red regions here sum
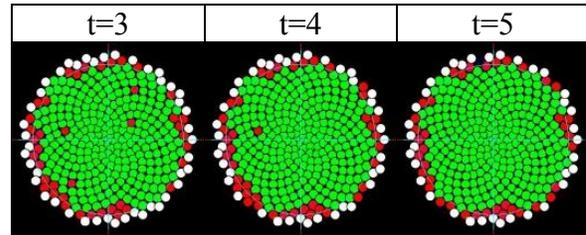
Illustration 26: Green filling-in for RF level 2

large lower level regions of red only patches. We shall discuss shortly how most of the filling-in is due to top-down flows, so this 2nd level needs to complete before the 1st level can, and explains why the lower level takes that much longer. There is also a one tick delay sending information down.

While level 2 completes its filling-in by t=5, one may assume that level 1 would complete by t=6. But that doesn't happen. This will be discussed shortly as the difference between interpolation and extrapolation which depends on RF alignment from one level to the next. What it also means is that even at this level, filling-in is not complete until t=14 from the lower levels. It is simply not apparent unless the statistics of the values are examined.

2.4.2.5 Green filling-in for RF level 3

One interior RF near the edge took 1 tick to fill-in at RF level 3 as shown in Illustration 27. It is questionable whether it should have been counted as interior. Again, this level starts displaying image information one tick after the lower level, at t=4.

t=4	t=5
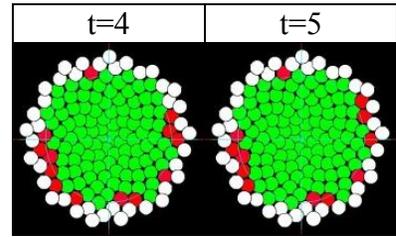
Illustration 27: Green filling-in for RF level 3

2.4.2.6 Blue filling-in for RF level 1

Blue cones make up 10% of this retina leading to many solitary islands of blue with vast blue free regions needing to be filled-in. While filling-in *stabilizes* within 14 ticks as shown in Illustration 28, it does not *complete*. This is unlike human perception which homogenizes fully. However, humans cannot prevent eye motion, and we hypothesize that filling-in within ECM would complete under ecological eye-motion. That hypothesis will be validated in Test 2.

t=2	t=3	t=4	t=5	t=6	t=7	t=8	t=9
t=10	t=11	t=12	t=13	t=14	t=15	t=16	
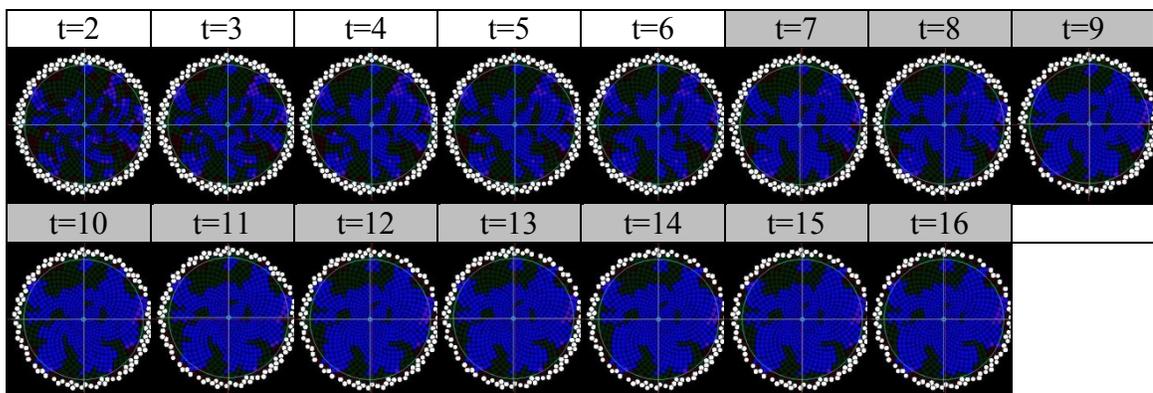
Illustration 28: Blue filling-in for RF level 1

Note that this filling-in can be characterized as a slow diffusion from blue-cones outwards to some imaginary border. Within the interior, it looks like diffusion inwards. However, because it never completes, this demonstrates that ECM does not have an

explicit diffusion mechanism. The emergence of diffusion and borders will be discussed later.

At time=7, this RF level shifts for the first time with a marked increase of filling-in (Illustration 29). In this case, the retina has jittered slightly downwards, and a little more to the left. Thus, the RF level has shifted a little bit to the right, into the
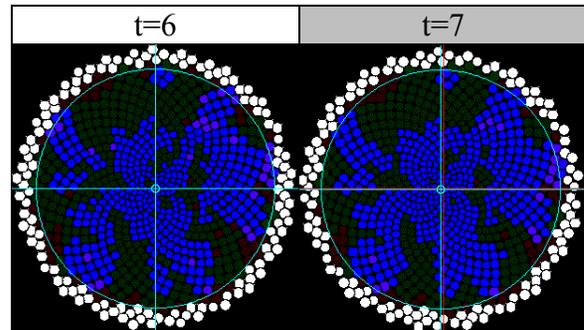
Illustration 29: Blue filling-in for RF level 1 (expanded view after initial jitter)

white headcentric borders, to compensate and *maintain information coherence*. It is this larger shift of information that facilitates the marked increase of filling-in. This will be demonstrated to greater effect in Test 2 where significant eye motion occurs. The headcentric RF borders are white indicating that they have not yet received any colour information. They cannot receive information from below because they are beyond the photoreceptors. They can only receive information laterally when it is shifted due to eye motion even as small as jitter, or from above.

2.4.2.7   Blue filling-in for RF level 2

The filling-in for level 2 stabilizes within 13 ticks but does complete (Illustration 30).

t=3	t=4	t=5	t=6	t=7	t=8	t=9	t=10

t=11	t=12	t=13	t=14	t=15	t=16		
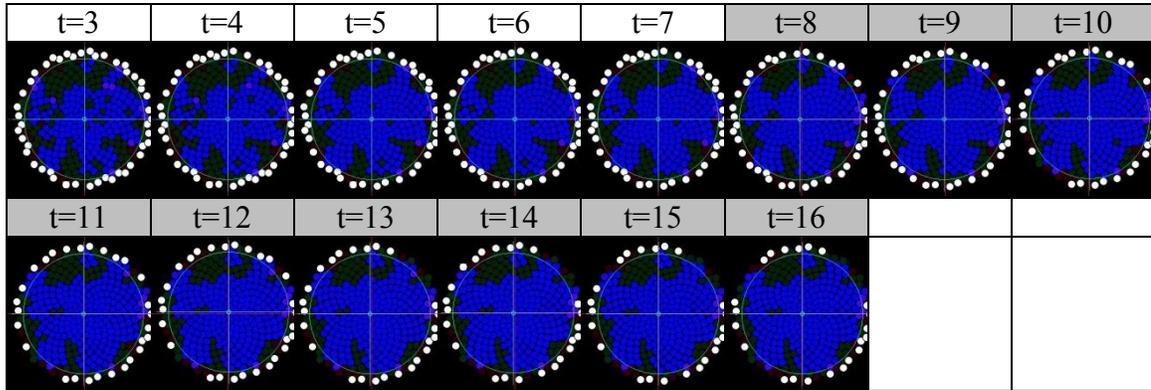
Illustration 30: Blue filling-in for RF level 2

Ordinarily, the higher level shifts in advance of the lower one but only if the lower one has undergone a shift. In this case, the lower one has not, so this level's shift first occurs at time=8 rather than 6. There is a marked increase of filling-in at this time (Illustration 31). In this case it may be difficult to discern if the interior point indicated by the yellow arrow was filled-in due to the shift, or due to bottom-up flows that were previously shifted. However, a look at level 1 indicates that the corresponding

t=7	t=8
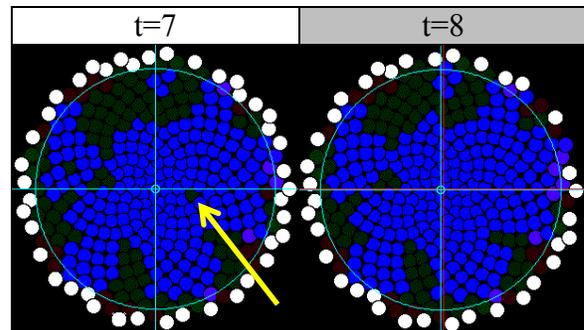
Illustration 31: Blue filling-in for RF level 2 (after initial jitter)

region has not yet filled-in, so the filling-in here is likely due to local lateral information flows (shifts). Of course, the software could have been instrumented to count which up, down or lateral flow of information was utilized. Indeed, once this level has filled-in, it can propagate the information downwards to help the lower level fill-in at this region. The path of filling-in can be convoluted, going upwards, downwards and sideways at will – there is no dedicated filling-in process.

2.4.2.8   Blue filling-in for RF level 3

The filling-in of level 3 RFs stabilizes in 12 ticks, but does not complete (Illustration 32).

t=4	t=5	t=6	t=7	t=8	t=9	t=10	t=11
t=12	t=13	t=14	t=15	t=16			
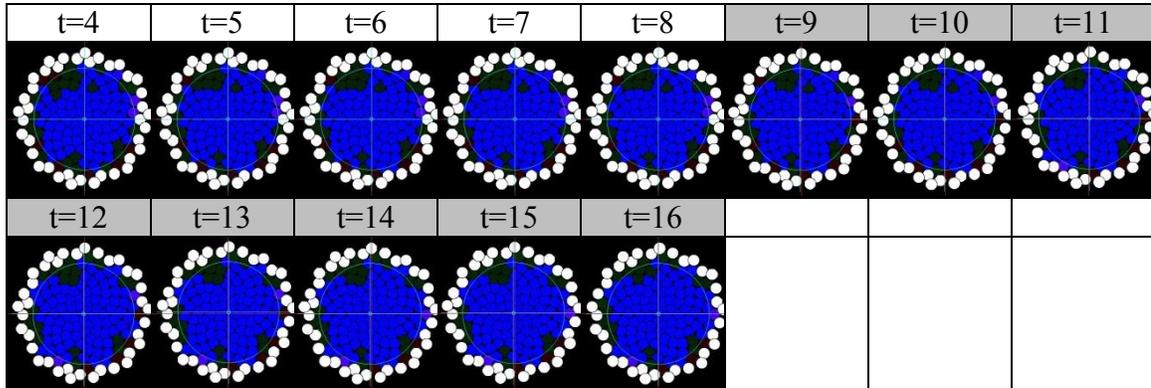
Illustration 32: Blue filling-in for RF level 3

There is additional filling-in when this level first shifts at time=9 (Illustration 33).

Illustration 33: Blue filling-in for RF level 3 (expanded view after initial jitter)

2.5   Test 2: Filling-in via eye motion

Filling-in of the blue patch did not complete, while it does for human participants. We hypothesised that in human participants completion occurs because of incessant eye motion. In particular, microsaccades occur at a minimum of once every two seconds (0.5Hz) and while not directly controllable, they are modulated by task and attention (Otero-Millan, Troncoso, Macknik, Serrano-Pedraza, & Martinez-Conde, 2008; Pastukhov & Braun, 2010). In this test, we simply validate our hypothesis. Further chapters will investigate eye movement in greater detail. Test 3 will characterize an alternative to eye motion.

2.5.1   Methods

The methods used are identical to those of Test 1 except as indicated herein.

2.5.1.1 Stimulus

Only the blue stimulus from Test 1 was used.

2.5.1.2 Procedure

Rather than seeing the stimulus at time=0, the agent was "instructed" to move its eye continuously to the right at 0.5′/tick. It takes ECM 7 ticks to plan and distribute eye movement information before the eye actually starts to move. At that point, the stimulus was made visible so that it can be viewed under constant linear motion. The stimuli remained in the field of view at all times.

2.5.2 Results and discussion

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s24). In this section, we summarize and discuss various highlights. We also extract filling-in times.

2.5.2.1 Photoreceptor mosaic motion constancy

As the stimulus is homogeneously blue, the photoreceptors show no difference in sensation as they move across the stimulus (Illustration 34). This can be contrasted with movement across a textured surface

Illustration 34: Photoreceptor mosaic response across blue stimulus

where the photoreceptors do show a difference, e.g., Illustration 42 (p52).

2.5.2.2 Blue-filling-in for shifting RF level 1

The filling-in for the level 1 RFs under motion stabilized in 9 ticks, but did not complete (Illustration 35). By tick 9, shifting has included all possible photoreceptors of the eye and further shifting would be superfluous. The image animation sequence shows that the behaviour of filling-in is best characterized as blue points leaving
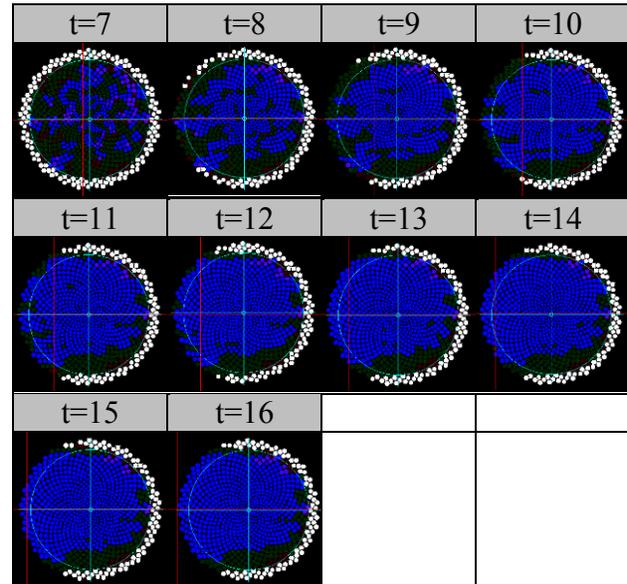
Illustration 35: Blue filling-in for shifting RF level 1

streaks as they shift to the left to compensate for eye motion to the right. If a blue point moves under an RF without any other blue information, that RF will use the laterally shifted blue information.

While filling-in did not complete (empty streak on the right), it was faster and closer to completion than without motion (Illustration 28). Counting the number of unfilled RFs could be a filling-in measure under speed. Had horizontal motion been reversed, the RFs on the right would have been filled-in. Had motion included a vertical component, the RFs at the top and bottom would have been filled-in. These will be demonstrated in the next chapter. Thus, we take this as confirmation of our hypothesis that eye movement would enable the complete filling-in of blue.

As it is, we note again that the bottom RFs acted as an imaginary border and filling-in did not diffuse across it. The headcentric RFs (white border) was filled in on the left as shifting was leftwards to compensate for eye motion rightwards.

2.5.2.3   Blue-filling-in for shifting RF level 2

The level 2 RFs were completely filled-in by 5 ticks as shown in Illustration 36.
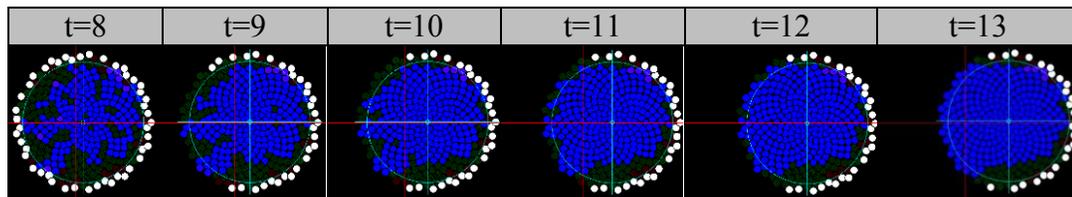
Illustration 36: Blue filling-in for shifting RF level 2

2.5.2.4   Blue-filling-in for shifting RF level 3

The level 3 RFs were completely filled-in by 2 ticks as shown in Illustration 37.
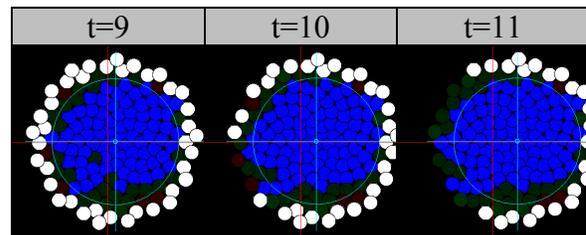
Illustration 37: Blue filling-in for shifting RF level 3

2.6   Test 3: Filling-in via larger RFs

In Test 2 we demonstrated that eye motion can facilitate filling-in. We also hypothesize that having larger RFs would do so as well, and we investigate this possibility in this test. While such larger RFs may not be biologically plausible throughout the visual system, they are indeed so over the blind spot, and this will be investigated in another chapter (p161). This section exists more to ensure complete characterization coverage of ECM, i.e., what are all the development parameters that could pertain.

2.6.1   Methods

All methodology is equivalent to the blue condition of Test 1 except as detailed herein.

2.6.1.1.1   Subject analysis summary

The development parameters for this particular agent were changed as detailed in the supplementary material (s2). These simply increase the size of its 2nd and 3rd level RFs while reducing their numbers to maintain similar fields of view.

2.6.2   Results and discussion

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s30). In this section, we summarize and discuss various highlights. We also extract filling-in times.

2.6.2.1   Blue filling-in for RF level 1

The filling-in of the interior level 1 RFs stabilizes by 9 ticks as shown in Illustration 38, although the exterior continues to fill-out due to larger higher level RFs.
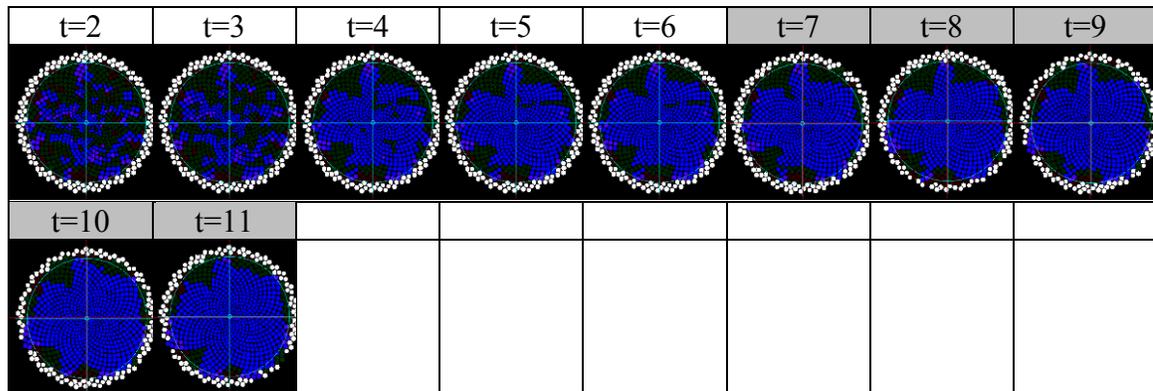
Illustration 38: Blue filling-in for RF level 1

Note that the RFs at this level are identical to those in Test 1, e.g., Illustration 28 (p42), but filling-in is faster and completes due to top-down flows. Moreover, the behaviour here can be characterized as diffusion all the way out to the extremities, while with smaller RFs, diffusion seems to be bounded by some invisible border – a phenomenon to be analyzed within the discussion section.

2.6.2.2   Blue filling-in for larger RF level 2

The larger RF level 2 interior is filled-in completely within 0 ticks, i.e.,

in parallel with and relative to the arrival of the first visual content one

tick later (Illustration 39). Thus it significantly helps level 1 fill-in (top-

down) at tick=4. Surprisingly, at tick=4, level 1 does not complete. That
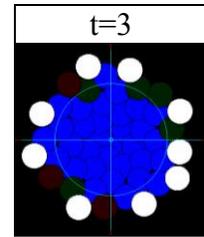
Illustration 39: Blue filling-in for larger RF level 2

is because while this RF displays its sampled values wherever those samples happen to be

spatially, the values maintain their statistical spatial properties. This contributes to lower

level *interpolation* but prevents *extrapolation* as will be discussed in the general

discussion section.

2.6.2.3   Blue filling-in for larger RF level 3

The larger RF level 3 interior is filled-in completely within 0 ticks

(Illustration 40).

2.7   Test 4: Finely textured patch with eye motion

In ECM, top-down filling-in is equivalent to spatial filtering which is
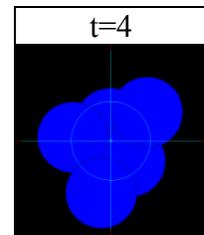
Illustration 40: Blue filling-in for larger RF level 3

known to cause blur by smoothing over fine spatial irregularities. Under modest eye

movement, lateral information flows are given precedence, and we hypothesize that these

do not suffer blur. In this test we validate this assumption so that comparisons to other

filtering accounts can be made. This test can be thought of as testing the counterfactual

hypothesis – that filling-in does not occur for textured surfaces. Indeed we use this test to

show that ECM does not perform isometric filling-in (there is no such explicit or

ontological mechanism to *explain* behaviour) but that "filling-in" emerges under

appropriate circumstance – it is an epistemic term for the *description* of behaviour (when

that behaviour occurs).

2.7.1 Methods

The methods used are identical to those of Test 1 except as indicated herein.

2.7.1.1 Stimulus

The stimulus as shown in Illustration 41 was used and is fully described in the supplementary material (s4). It consists of a fine texture of red dots over a green background on the left, and the opposite on the right.
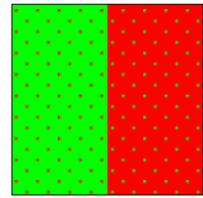
Illustration 41: Fine texture stimulus

2.7.1.2 Procedure

The virtual agent was positioned such that the size of the cone matched the size of the textured points near the middle of its mosaic. The virtual agent was situated at 180cm from the computer screen. It was instructed to move its eye upwards at $0.5'$/tick in a continuous path once the stimulus was presented. The stimuli remained in the field of view at all times.

2.7.2 Results and discussion

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s36). In this section, we summarize and discuss various highlights. We also extract filling-in times.

2.7.2.1 Photoreceptor mosaic motion

Because the stimulus is textured, eye motion should be visible by looking at the photoreceptors as in Illustration 42, and unlike the apparent motion constancy of Illustration 34 (p46). If the eye is moving up, the stimulus will appear to move down. This is harder to see in the frame-by-frame sequence as shown.

t=6	t=7	t=8	t=9
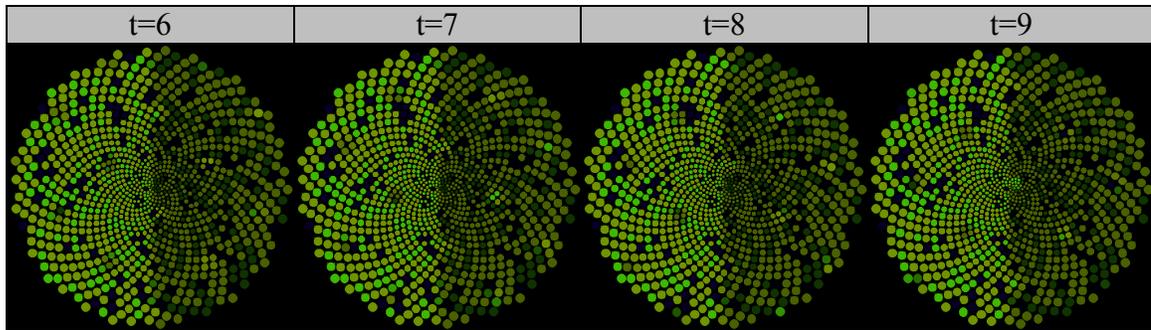
Illustration 42: Photoreceptor mosaic under motion

2.7.2.2 Photoreceptor mosaic difference showing motion

The motion can be made more salient by computing the absolute value of the difference (see page 394) between the current value and the previous one, as shown in Illustration 43.

t=6	t=7	t=8	t=9
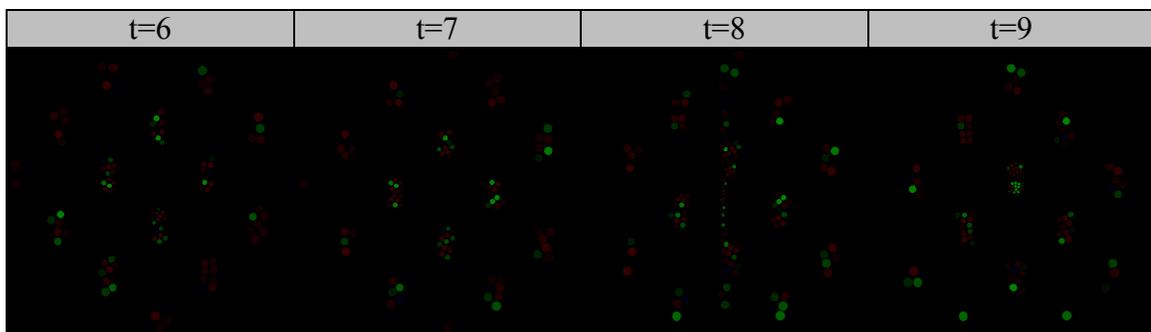
Illustration 43: Photoreceptor mosaic differences under motion

The points are elongated as there would be a negative difference where the old point disappears and a positive difference where the new point reappears.

2.7.2.3 Texture dynamics at RF level 1

Under motion, the level 1 RFs do not fill-in as shown in Illustration 44 and in greater detail within Illustration 45.

t=2	t=3	t=4	t=5	t=6	t=7	t=8	t=9

t=10	t=11	t=12	t=13	t=14	t=15	t=16	
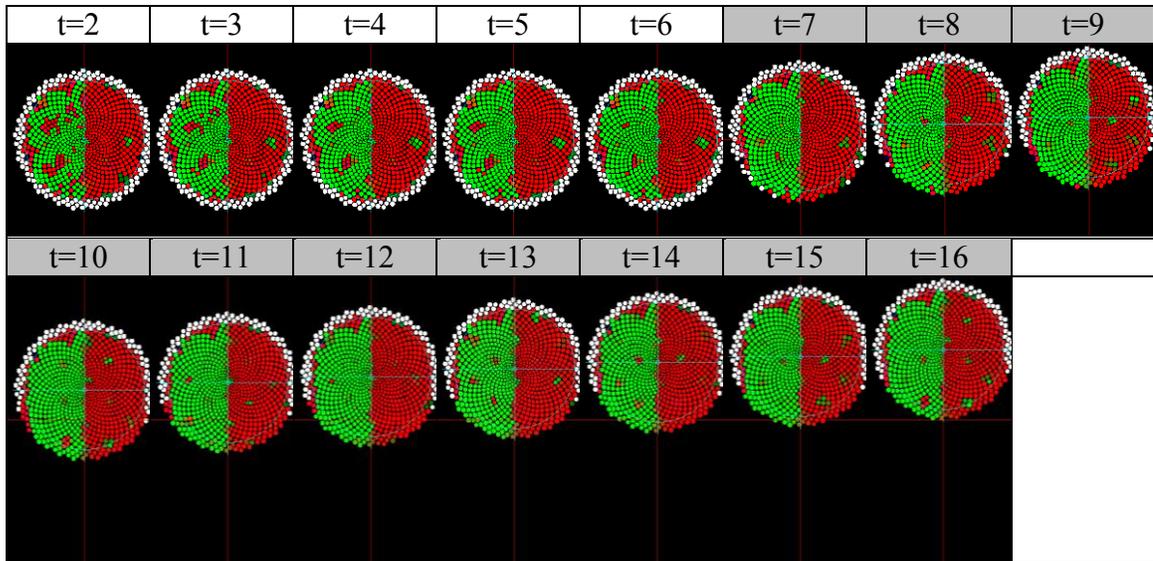
Illustration 44: No RF level 1 filling-in of textures under motion

Note that if an RF covers half a green and half a red cone, then it would be shown as brown, but the texture points are smaller and less salient when mixed. This is not to be taken as a form of filling-in. Indeed, it does highlight a problem of showing information in colour. Each RF maintains the LMS information separately, so ideally each could be shown individually necessitating three times as much visual real estate.

Before motion starts, and even a few times afterwards, it does look like some of the green points on the right have been filled-in with red, but that's because it is hard to distinguish red from brown. When expanded, frames 12 and 13 show that filling-in of textured points does not complete as in Illustration 45. This should be compared to Illustration 25 (p40) which has no brown artefacts.
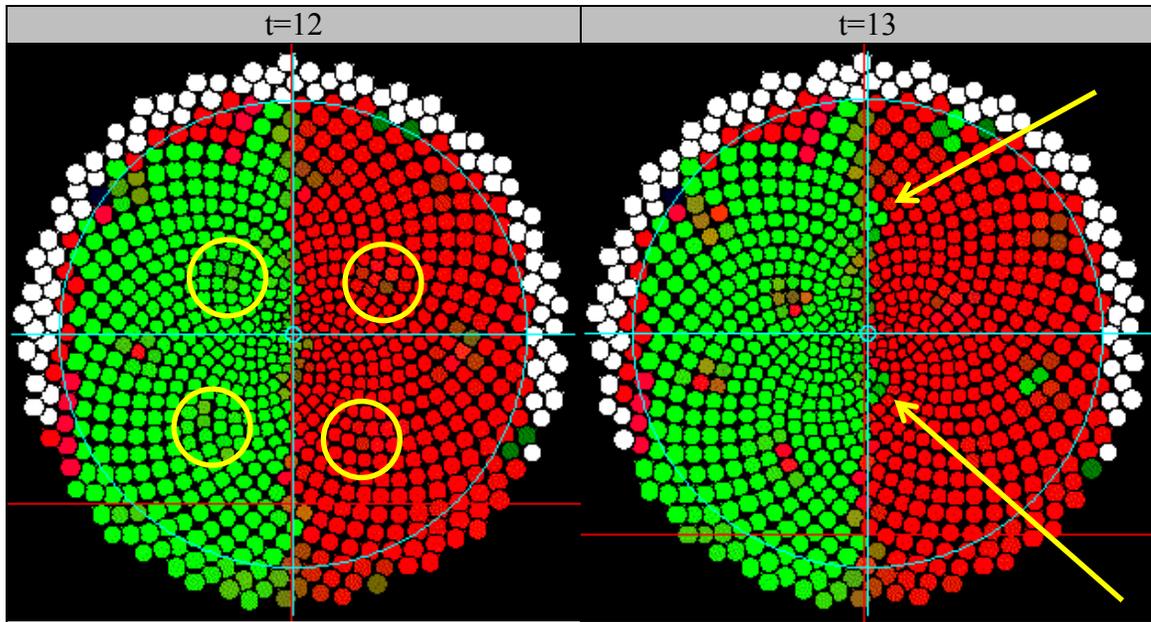
Illustration 45: No filling-in (brown points) for RF level 1 (expanded view)

Compare our no filling-in account with the Grossberg based account of surface filling-in that incorrectly eliminates salient noise from perception (Illustration 12, p22). Also note that while filters in general blur information, the edges between the square and background are straightened via contour completion of the boundary contour system (BCS). In our figure, the imperfect borders between the green and red patch are broken up by texture artefacts, clearly indicated on frame t=13. These are salient and their irregularities are not homogenized at our sensory level of processing.

Our account is based on sampling statistics. If an L, M, or S value is sampled, it is never replaced by a homogenization process. Upon eye movement, it can be used laterally to fill-in an adjacent RF only if that RF is missing such sampling information, and in this case, there would be no loss of resolution, assuming motion was handled accurately. If lateral information is still not available, then information from above can be used, but in this case it is an average – an interpolation (but not extrapolation) from a larger RF.

Nevertheless, even in a grey-scale approach, one white textured point sample in an RF consisting of a mostly black patch will statistically be shown as fairly dark.

2.7.2.4   Texture dynamics at RF level 2

Even at RF level 2, there is no filling-in as shown in Illustration 46. In some cases, the lack of filling-in is not salient unless the images are examined under greater details, such as done in Illustration 45.
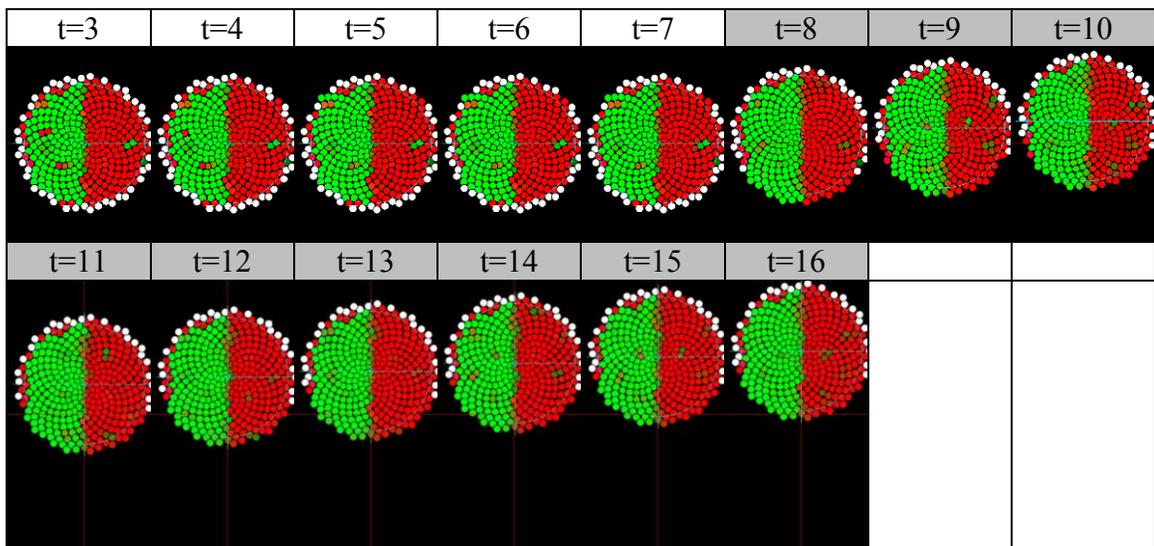
Illustration 46: No texture filling-in for RF level 2

2.7.2.5   Texture dynamics at RF level 3

Even at level 3, where the texture point is less that 50% of the RF size, there is still no filling-in as shown in Illustration 47. The lack of filling-in is much more salient in the animated images. There is no spatial homogenization that eliminates statistical sampling.

t=4	t=5	t=6	t=7	t=8	t=9	t=10	t=11

t=12	t=13	t=14	t=15	t=16			
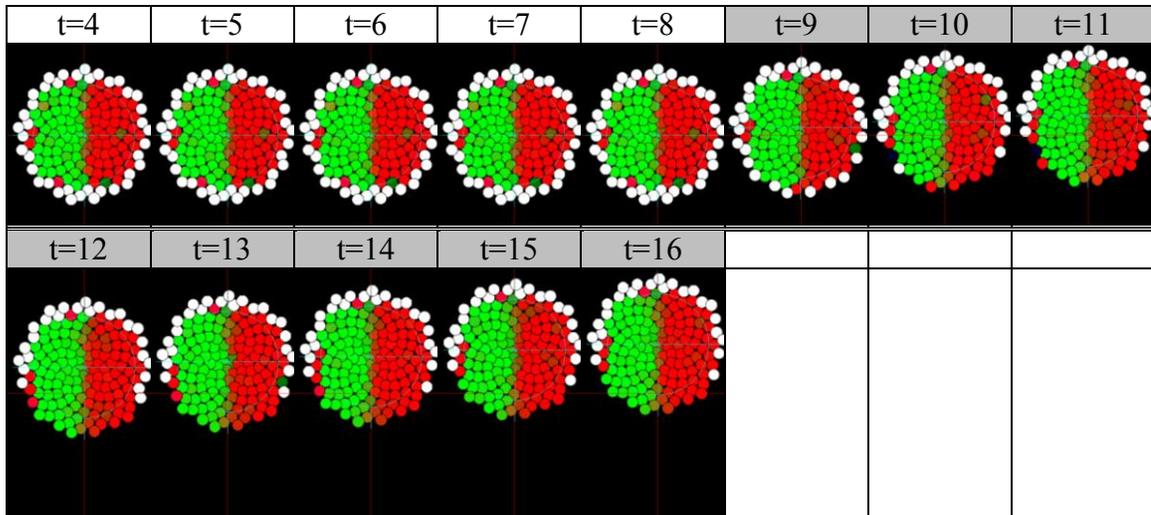
Illustration 47: No texture filling-in for RF level 3

2.8   General results and discussion

In this section we discuss the behavioural emergence of filling-in, diffusion and borders

via an explanation for why ECM interpolates but does not extrapolate.

2.8.1   ECM functional characterization

In most cases, the initial filling-in stabilizes before eye movement begins. This implicates

the *handle missing data* function using top-down flows of information. When jitter starts,

another burst of filling-in occurs and stabilizes. This implicates both the *maintain*

information coherence function responsible for shifting in conjunction with the *handle*

missing-data function. In this case, most of the information comes in laterally, but

sometimes downwards. In actual fact, the source of filling-in information can flow

around the system, much like Brownian motion, moving through function to function,

and level to level before being used to fill-in an irregularity.

With larger eye movements, most of the filling-in information comes laterally from the

maintain information coherence function. Nevertheless, this only shifts the information

laterally, and the *handle missing data* function actually performs the ultimate filling-in.

56

This typical scenario is further characterized in a subsequent chapter. These two functions interact (or collaborate) at the locus of an RF.

2.8.2   Interpolation but not extrapolation

To explain the ECM difference between interpolation and extrapolation, we will refer to the one dimensional simplification of the S (Blue) cone channel as shown in Illustration 48. Each receptive field has a physical visual extent. So the highest receptive field shown
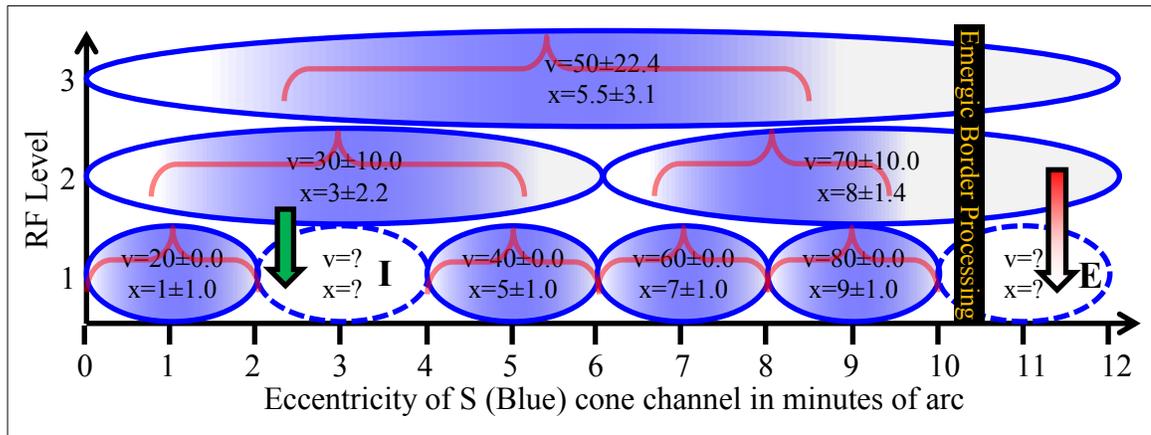
Illustration 48: Interpolation vs. Extrapolation
Showing the mean of each RF value ± 1 standard deviation.

has a center at 6′ and a diameter of 12′. However, all values flowing within ECM are encoded with statistics over both their content and spatial characteristics. In this case, the content is the amount of S cone activation representing Blue. Each level 1 RF is assumed to have two point samples representing their spatial extremities, 0′ and 2′ for the leftmost RF. If the sampling of blue cones is spatially summed up the hierarchy, the statistics are as indicated in the illustration, with ±ranges representing the standard deviations. The interior shading represents these spatial statistics, as well as by {braces}.

In this example, the RF marked **I** is amenable to interpolation. Assuming no eye motion, it will ignore all lateral values as they are out of its physical extent. As it overlaps the

57

value coming from above 100% (green arrow ↓), where overlap value borders are defined by their standard deviation, it will take the value as is, namely v=30±10.0 and x=3′±2.2. The situation with the RF marked **E** (for extrapolation) is different. Its physical spatial extent is beyond the spatial statistic of lateral and top-down values (red arrow ↓). So while receiving them, it will ignore the values completely. If there was some overlap, it will take the value as given weighted by the percentage of spatial overlap. Note that the RF marked **I** will also reject lateral values as we assumed there was no jitter and therefore these lateral values from surrounding RFs do not overlap with the special extent of the RF **I**.

Thus, RFs to be interpolated will appear to fill-in, while those on the extremities will be bounded by spatial extent causing border processing to emerge. The RF may appear to fill-out slightly, due to some lateral overlap with jitter, but the filling-out is bounded by spatial statistics.

While spatial statistics are critical to the routing of information within a locally connected neural region, they also allow for spatiotemporal integration that could be used for hyperacuity tasks.

2.8.3   Sensory irregularities vs. image irregularities

Most alternative visual models do not distinguish between *sensory* irregularities and *stimuli* or *image* irregularities. That is why they will fill-in for all cases, and will blur fine texture. Our account is explicitly not about filling-in. All sensory information is treated as samples, and their irregularities become inconsequential if their spatial positions are

taken into consideration. We do not need to explicitly regularize sensory samples, as this is perfectly accounted for in the dynamics of eye motion under a statistical encoding regime.

3   Chapter: An emergic model of filling-out for trans-saccadic integration

The eyes undergo incessant motion yet the world appears stable. Moreover, while each eye fixation samples a subset of the world, we perceive the larger *headcentric* or *spatiotopic* view as if all the smaller *retinocentric* snapshots had been stitched together. This collective behaviour is called *trans-saccadic* perception across saccades (Melcher & Colby, 2008) and peri-saccadic perception near and during saccades (Hamker, Zirnsak, Calow, & Lappe, 2008). We call it *filling-out* as that relates directly to the mechanism we hypothesize and demonstrate in this chapter. Regardless of mechanism, all previously proposed models, e.g., (Feldman, 1985), seem to require some form of 1) coordinate transformation and 2) visual memory to underlie stable *perception* "of the large" so that old *sensory* information about one place is maintained while the eye is 3) accumulating new evidence elsewhere. The only exceptions are models claiming that image detail and stability are illusions.

There are numerous conjectures as to how image stability and integration come about, but few computational models. One reason for the paucity of models is due to complexity – the final model would need to integrate the sensorimotor system of the eye, with possible coordinate transformations systems, a short term visual memory system, and a data fusion system. Many models attack only a small subset of this complexity. We suspect that the actual reason for model scarcity is that many of the basic conceptions are problematic, ill-defined or overly simplistic. For example, most receptive fields (RFs) are retinocentric, but can be modulated by gaze (Andersen & Mountcastle, 1983) implying an implicit

headcentric encoding (Hamker & Ziesche, 2011). Other retinocentric RFs are known to *shift* under various cases (Kusunoki & Goldberg, 2003; Womelsdorf, Anton-Erxleben, Pieper, & Treue, 2006; Womelsdorf, Anton-Erxleben, & Treue, 2008) or *remap* (Merriam et al., 2007). The phenomenon of RF shifting, and gaze modulation are not compatible with the concept of *retinocentric*, nor of "classical" *receptive field* in general (Connor, 2001). Perhaps for this reason, the coordinate system for some brain regions is still in doubt (Olshausen & Anderson, 2010). While an eye position invariant spatiotopic or headcentric map is thought to be necessary for image stability, some continuous remapping schemes can make this superfluous (Colby, Duhamel, & Goldberg, 1995). Indeed, a headcentric map amounts to a Cartesian theatre watched by the conscious homunculus (Dennett & Kinsbourne, 1992; Dennett, 1991) – it cannot explain perception, only infinitely regress and pass off the "hard" problem. Thus, the concept of image stability is a phenomenological supposition bringing the problem of identifying the locus of consciousness to the fore – technically, image stability is within that locus.

Visual memory that facilitates motion perception is also problematic. For example, the *persistence of vision* is considered a myth (J. Anderson & Anderson, 1993) plagued by definitional problems (Long, 1982), yet some form of memory seem to be required for handling integration, motion and stability, and thus many kinds, and interactions therein, have been proposed, e.g., sensory, iconic and visual short-term memory (VSTM). However, it is not at all obvious what form memory ought to take nor how it ought to be integrated within a system such that it would not lead to artefacts associated with persistence such as smearing or ghosting (Öğmen & Herzog, 2010), but would retain

peri-saccadic phenomena such as spatial compression (J. Ross, Morrone, Goldberg, & Burr, 2001). Finally, saccadic suppression studies show that much of visual process during a saccade is reduced, while change blindness studies show that few details survive a saccade (and vision in general). Taken together, these studies suggest that the phenomena of peri-saccadic and trans-saccadic perception are merely filling-in illusions. The world itself is the source of detail, and we recreate detail anew every time – there is no stabilizing or integrative phenomena needing to be explained. Of course, the phenomena of change blindness and suppression ended up being much more complex and context dependent, so much so, that many of the original conclusions have been found to be overly strong and illegitimate (Simons & Rensink, 2005) leading to numerous misconceptions, pitfalls and errors of interpretation (Jensen, Yao, Street, & Simons, 2011). These phenomena resist oversimplification and require a balanced view (Castet, 2010). Unfortunately, a more unified or holistic approach is more difficult to model.

Since the earliest days of change blindness research, image detail, if any, was consigned to a short-term low-level *sensory* system with change blindness occurring at long-term higher *perceptual* levels. In this chapter we present a low-level sensory model of filling-out for peri-saccadic and trans-saccadic integration which supports the phenomena of image stability and visual memory. The model is based on the Emergic Cognitive Model (ECM) architecture (p323). ECM can clarify conceptions precisely because it does not prematurely reify them. For example, ECM does not have any specific mechanism to fill-out, stabilize nor to transform or remap one coordinate system to another. However, we hypothesize that both of its basic functions – *handling missing data*, and *maintaining*

information coherence – could interact with test stimuli and measurement to cause this filling-out, image fusion, image stabilization and coordinate transformation phenomena to emerge. In other words, rather than model the phenomena directly, we model other functional aspects explicitly, and the phenomena in question emerge indirectly. As part of characterizing the relative contributions of the basic functions to the main phenomena, we also discovered that out of the model emerged a dynamic *infocentric* coordinate system which offers an illuminating clarification to the retinocentric and headcentric debates for image stability, and offers a unique hypothesis as to the locus of consciousness. The phenomenon of visual memory and fading of those memories also emerged. As before, our results are demonstrated visually and in colour.

3.1 Introduction

In this section we review the phenomenology and models associated with maintaining visual stability across eye motion, and the integration of visual information from different locations in space. Much of this section also serves as background material to this chapter and subsequent ones, as this is the where we first deal seriously with the problem of motion – the previous chapter dealt mostly with static (albeit jittering) eyes.

3.1.1 Eye motion

The eye is under incessant motion. This is readily demonstrated in eye position traces such as Illustration 49. Each point/vertex represents a short period of fixation where a detailed conceptual *snapshot* is taken. However, even a voluntarily focused and "steady" fixation at one point, e.g., Illustration 50, cannot
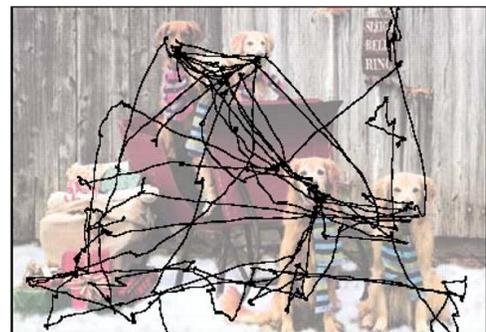
Illustration 49: Incessant eye motion. From(Otero-Millan et al., 2008).

prevent the eye from continuously jittering up to 6″ at 40-100 Hz (Eizenman, Hallett, & Frecker, 1985), nor of taking involuntary micro-saccades away from the fixation point typically up to 12′ every 0.5 Hz (Martinez-Conde, Macknik, Troncoso, & Hubel, 2009; Steinman, Haddad, Skavenski, & Wyman, 1973). People are not aware of either involuntary eye movement type (jitter or micro-saccade) as the world remains stable. They are also not aware of the change of detail between the fovea and the periphery – the entire picture looks detailed even if the resolution was vastly reduced in the periphery where fixation did not occur. This can
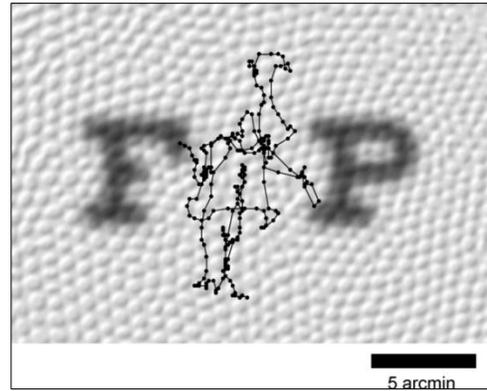
Illustration 50: Fixational eye movement.
From (Pitkow, Sompolinsky, & Meister, 2007), superimposed on a foveal cone mosaic from (Roorda & Williams, 1999).

be demonstrated via gaze-contingent displays (Loschky, McConkie, Yang, & Miller, 2005) that only show detail where the eye is currently watching. The pan-field colour effect (Illustration 15 on page 27) also demonstrates that large amounts of missing colour detail are not noticed. Detail constancy appears to be an illusion.

3.1.2 Trans-saccadic perception

Each fixational snapshot of the world is conceptually fused together as shown in Illustration 51. This descriptive metaphor is known as trans-saccadic perception (TSP), but it is inherently implausible as stated (Melcher & Colby, 2008) – the visual system does not take individual pictures, the "pictures" are only

Illustration 51: Trans-saccadic perception.
From (Melcher & Colby, 2008) Fig. 1: Fusion of separate snapshots. Collage by (Pozzi, 2001).

detailed in their centers, integration does not happen after the fact with a set of pictures, and the brain does not have a stitching system that seems to require a point-to-point connection from each retinocentric coordinate (of each snapshot) to each headcentric coordinate for the entire collage. Nevertheless, if the global canvas serves as the stable image reference point, this metaphor forms a starting point for investigating actual mechanisms of image stability whereby the larger image continually accumulates additional visual information. A problem with such folk-theoretic notions is that they can limit scientific inquiry and theorizing, while a computational account can clarify.

3.1.2.1 Filling-out terminology

If one considers the center of the larger canvas as the headcentric fixation point, then the involuntary microsaccades that move the eye around serve to *fill-out* that panorama with extra detail, hence giving our name for the phenomenon. Note that *filling-out* is our term for the phenomenon as it is the best description for the behaviour of our model as shown in the model test sections, and for understanding our computational mechanism. Such renaming of phenomena is common – Troxler's *fading* was renamed to *filling-in* as this was closer to a new interpretation and description of mechanism (from a "surface adaptation" explanation to a "diffusion from an edge across a surface" explanation).

Of course, when saccading rightwards, if the old information from the retinal region filled-*out* the periphery on the left, for example, then it is equally possible that information from the right periphery filled-*in* the retinal region. Ordinarily, this filling-in would not occur as the retina is already receiving stimuli directly, however, during an eye-blink (p263) or saccadic suppression, such filling-in could explain the peripheral preview effect (Khayat, Spekreijse, & Roelfsema, 2004).

3.1.2.2   Change blindness

Changes are often hard to detect during saccades and this specific case will be discussed shortly. However, it should be noted that change blindness can be induced in all sorts of contexts (Rensink, 2002), including when there is no eye motion and when dynamic visual stimuli, such as an extremely salient gorilla thumping its chest, move across the fovea and across the spatial extent of focused attention (Simons & Chabris, 1999). See Illustration 20 on page 33.

The various change blindness phenomena have affected the modelling of trans-saccadic perception in two ways. Firstly, they can suggest that the entire phenomenon of detail constancy and stability is a typical filling-in illusion that does not need any explanation. Fusion could simply be at a higher abstract level (Irwin, 1991). An active vision model would simply acquire detail only when needed. Our model and alternatives would, therefore, be superfluous! In particular, there would be no need to maintain a spatiotopic visual buffer that would accumulate information (Rensink, 2002). Secondly, it can suggest a nature of representation to be used in modelling, for example whether it is non-representational, symbolic, or a hybrid based on proto-objects.

We should point out that ours is a sensory account and whether our accumulative visual memory reaches consciousness is a separate affair. Our account, for example, will allow for detailed sensory information to make its way into the blind spot where, surprisingly, there does exist receptive fields and where perception can be directed. This is detailed in a latter chapter (p161).

3.1.3 Peri-saccadic perception

Intra-saccadic perception occurs during the saccade, while peri-saccadic perception occurs near and during the saccade as it includes some pre- and post-saccadic phenomena. Space compression is an example (J. Ross, Morrone, & Burr, 1997), although it is hard to distinguish between attentional causes (Liverence & Scholl, 2011) or saccadic ones (Binda, Cicchini, Burr, & Morrone, 2009). Regardless of when eye motion occurs, the phenomenon of image stability remains as does the integration of new information with old. Our focus remains on these last two effects, while many other models focus on specific peri-saccadic effects.

The following additional peri-saccadic phenomena are worthy of note as they relate to possible image stabilization mechanisms.

3.1.3.1 Corollary discharge

In advance of a saccade, the brain emits a corollary discharge (Sommer & Wurtz, 2008). This could help with image stability in two ways. If the corollary discharge represents an efference copy of the saccadic motor commands, it could be used to adjust frames of reference proactively rather than reactively to achieve image stability. The advanced

shifting of receptive fields is one such consequence. The corollary discharge could also be used to suppress visual information during a saccade to prevent blurry images from reaching consciousness.

3.1.3.2 Efference copy

There are unresolved and ongoing debates on how much of eye motor control and perceptual stabilization is due to 1) advanced, endogenous, planned, non-retinal motor *efference* outflow
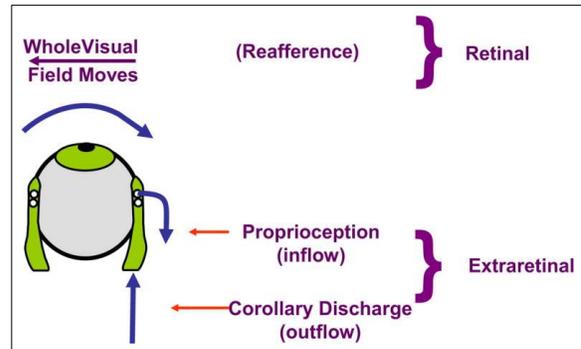
Illustration 52: Signals available for achieving visual stability.
From (Wurtz, 2008), Fig. 2

feedforward signals; 2) proprioceptive, non-retinal motor *reafference* inflow feedback signals; and 3) exogenous retinal visual stimuli *exafference* feedback signals. Likely all affect perception and stability to various degrees (Wurtz, 2008). Even this terminology is disparately used as shown in Illustration 52 – if an object in the environment moves, its motion relative to the world equals its motion as seen by the eye, minus the motion due to the eye, i.e., exafference – reafference. The reafference is simply the efference copy, possibly adjusted to account for the difference between planned and actual eye motion.

The corollary discharge can represent an advanced feedforward efference copy of a planned saccade which can be used in a compensatory manner to achieve spatial constancy. However an efference copy is still problematic (Bridgeman, 2007). For example, it represents a planned move rather than an actual move, and therefore eye drift of up to 30′ must be taken into account, possibly by using proprioceptive feedback (a sensory *reafference* signal). Ordinarily, such a feedback control system would ensure that

drift phenomenon would not occur in the first place, but it does indicate a possible functional role (Donner & Hemilä, 2007) such as sharpening blur (Hammett, Georgeson, & Gorea, 1998). It is possible to show that the visual stimuli itself (exafference signal) is used for much of fixation (Poletti, Listorti, & Rucci, 2010) and possibly drift related corrective saccades (Kagan, Gur, & Snodderly, 2008). In any case, during a correction, or microsaccadic flick (St Cyr & Fender, 1969), there is still a need to maintain spatial constancy (Cullen, 2004), and it is possible that a corollary discharge/efference signal/extra-retinal information is used to supress coordinate transformations during such corrections. Alternatively, corrections can somehow be masked due to retinal stimuli changes.

The spatial compensation mechanism still needs to be specified, and the advanced shifting of RFs could be the neurobiological mechanism which underlies the phenomenon of perceptual stability.

3.1.3.3   Saccadic suppression

One cannot see one's own eyes move in a mirror (Dodge, 1900), ostensibly to prevent image smear or blur. This test is so trivial and the suppressive effects so absolute, that this was the main reason (along with general change blindness) that peri-saccadic research and modelling was ignored – there was nothing to see or do (Castet, 2010). However, even in 1900, Dodge refuted the momentary visual anæsthesia hypothesis because bright lights would leave streaks, and the holes on a rotating disk of holes (with light shining through) could only be resolved distinctly when the eyes moved rather than appear fused when the eyes where static. Even moving print could be perceived during

eye motion, albeit shadowy. Dodge preferred the term *inhibition,* and the danger of the term *suppression* in connoting oversimplified absolutes is reiterated to this day (Castet, 2010).

There is an unresolved and ongoing debate on how much of image stability is due to 1) advanced, *active*, non-retinal saccadic *suppression*, e.g., signalled via a corollary discharge; or 2) *passive*, retinal stimuli dependent *masking*. Likely both affect perception and stability to various degrees (Wurtz, 2008). In Test 5, we demonstrate a masking phenomenon.

While it may be obvious how suppression or masking reduce image blur, it is not clear how these mechanisms relate to image stability, the phenomenon we are actually trying to explain.

3.1.3.4   Receptive field shifting

Prior to saccades, many receptive fields are known to shift in what is known as predictive remapping (Duhamel, Colby, & Goldberg, 1992; Hall & Colby, 2011; Merriam et al., 2007). While the terms *shifting* or *remapping* can suggest a mechanism that would transform a retinocentric coordinate based cortical map, to a visually stable headcentric map, this predictive remapping is in the wrong direction for a headcentric transformation (Burr & Morrone, 2011). No spatiotopic map has been found (Wurtz, 2008).

It is currently thought that such shifting enables a transient retinotopic map that is updated on each saccade (Wurtz, 2008). Why this leads to stability is often not

mentioned. Indeed, taken to extreme, the update time is reduced to zero if each saccade is infinitesimally small, i.e., the transient retinotopic map becomes a continuously perfect retinotopic map. The update mechanism does not buy anything towards stability.

One trouble with the shifting account is its transient nature. By the end of the saccade, the RF must be shifted back to maintain retinocentricity. It is not clear that there is a signal for this. However, in a double-step saccade task where the 2nd target disappeared before the first saccade occurred (Hallett & Lightstone, 1976), there is a requirement to remap the 2nd target across both saccades (Medendorp, 2011). This implies a less than transient visual memory, which has not yet been found.

Similarly, when the eyes of monkeys had been electrically stimulated prior to a saccade, eye movement was still accurate, leading Mays & Sparks (1980) to suggest that saccades are encoded in spatiotopic target coordinates and not by relative distances from last foveal position. Of course, this implies that the attentional sensorimotor system driving the generation of saccades must be functionally spatiotopic as well.

3.1.4 Filling-out is an epiphenomenon?

It is possible that any specific filling-in or filling-out phenomenon is an illusion and merely represents a particular sensorimotor engagement with the environment (O'Regan & Noë, 2001). So in the next few sections, we will provide examples of other filling-out phenomena that may bear on our primary one, and on supporting neurobiological evidence.

3.1.4.1 Boundary extension

Boundary extension, as demonstrated in Illustration 53, is a filling-out phenomena whereby an image is recalled, recognized or drawn larger than originally presented with pronounced extension or completion of object boundaries (Intraub & Richardson, 1989; Intraub, 2012). It can occur across mask intervals as short as 42ms and may therefore be involved in trans-saccadic spatial layout (Dickinson & Intraub, 2008). Unlike the primary

Illustration 53: Boundary extension.
From (Intraub, 1997; Intraub & Dickinson, 2008; Intraub & Richardson, 1989).

phenomenon, it must operate at a higher level of grounded object and schema perception (Hubbard, Hutchison, & Courtney, 2010). In this case it is a true filling-out *extrapolation* beyond the edges of a view (Park, Intraub, Yi, Widders, & Chun, 2007).

This extrapolation need not occur at image encoding. All objects seem to have a canonical size (Konkle & Oliva, 2011). The gist of a scene is encoded quickly from the visual features and then these visual features are discarded leaving participants blind to their changes (Sampanes et al., 2008). Thus, a scene may be encoded in a compact form, and its recreation with extended boundaries simply reflects the fact that the original boundaries or differences from canonical form included details that were not encoded.

This again illustrates that detail constancy could be an illusion not actually captured even in short term memory.

Similarly, by investigating after-effects Burr & Morrone (2005) have shown that low-level contrast after-effects did not transfer across saccades while high-level descriptions such as orientation, form, and especially face representation did transfer across to spatially identical, but not retinally distinct locations. This implies a buildup of high-level (gist like) information across space and the discarding of low-level details.

There are no computational models associated with boundary extension.

3.1.4.2   Neurobiological support

The phenomenon of image stability does not lend itself to mechanistic intuitions until a mechanism for conscious perception is understood. Thus supposed stabilizing and filling-out functions must be considered illusive until supported by biological structure. Here we present a few neurobiological underpinnings.

3.1.4.2.1   Headcentric receptive fields

The existence of headcentric or spatiotopic receptive fields can imply that if the eye moves rightwards then there is now some headcentric information on the left that may need to be filled-out as it is no longer receiving retinal input. While the spatiotopic status of RFs is under continuous debate (Gardner, Merriam, Movshon, & Heeger, 2008; Golomb & Kanwisher, 2011), there exist quite a few areas that are putatively headcentric (Burr & Morrone, 2011; Crespi et al., 2011).

3.1.4.2.2   Beyond the view

While headcentric RFs can underpin image stability, by themselves, they do not require any *extra-retinal* filling-out, i.e., beyond the reach of the photoreceptors. The headcentric field of view can simply be smaller than the retinotopic one. Any visual information within the headcentric periphery that moves out of view could be discarded. Indeed, the periphery already has less spatial and colour detail for which we are ordinarily unaware. Here we provide neurobiological evidence that the headcentric field is larger than the retinotopic one implying a need to fill it out.

The phenomenon of apparent motion has been found to operate with such extra-retinal information (Szinte & Cavanagh, 2009, 2012). There are "some circumstances" where afterimages can be moved into impossible locations beyond the view, and can still be perceived (Hayhoe & Williams, 1984).

Visuospatial information is integrated across the senses, and saccadic activity can be cued by auditory stimuli. Tark & Curtis (2009) found that FEF activity can be maintained even for locations behind the head or "off the map". In other words, the FEF can encode saccade related information in retinal and extra-retinal space. This does not mean that all visual processing areas are larger than that implied by their retinal field of view. However, behavioural evidence suggests that much of the visual system may be so involved.

Perceptual-motor research can suggest numerous views beyond the current reach of the photoreceptors (Previc, Declerck, & De Brabander, 2005). Some are 360° as shown in Illustration 54. Neural activity representing such unseen space has been found in the parietal cortex (Schindler & Bartels, 2013).
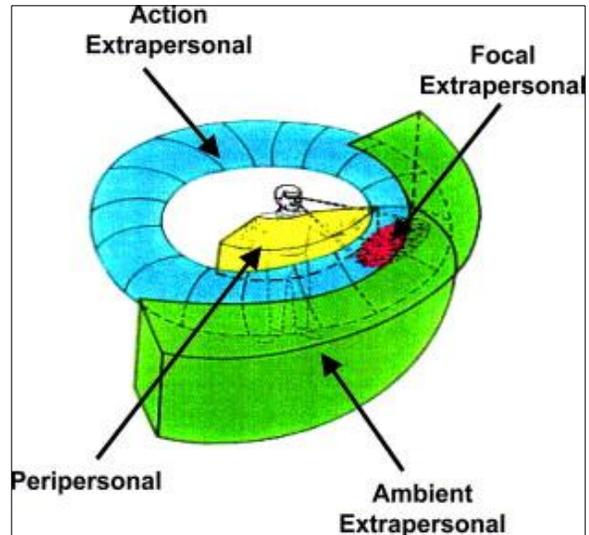
Illustration 54: Perceptual-motor 3-D space
From (Previc et al., 2005), Fig. 1.

3.1.4.3 Temporal extended spatial memory

When performing a task, it seems that the visual system performs 'look-ahead' fixations that may result in using an object beyond the current focus of attention. This implies that some information must be retained beyond the view (Tatler & Land, 2011).

3.1.4.4 Filling-in outwardly

The term *filling-out* is not as common as *filling-in*, and both can characterize a multitude of phenomena. In this section we review the most commonly assumed filling-out phenomena simply to distinguish it from ours and to clarify our scope.
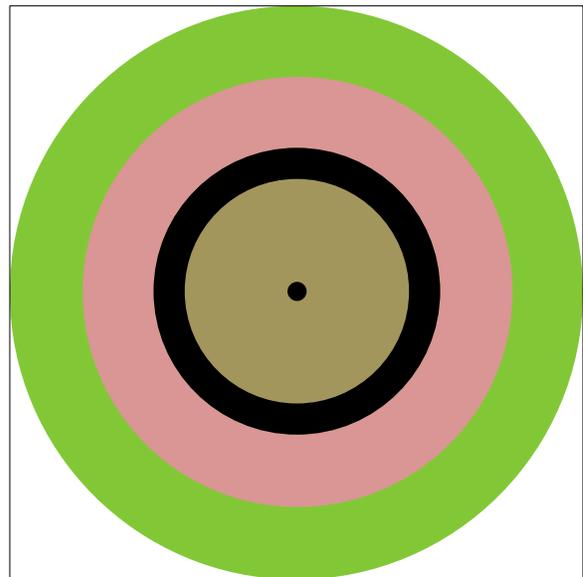
Illustration 55: Filling-out direction of filling-in.
After (Hamburger et al., 2006) Fig. 9: Black inducer filled inwardly and often outwardly.

The standard account of filling-in involves extended fixation at the center of a surface which leads to the fading of that surface (Troxler, 1804) and a filling-in from the inducing surround. This direction of filling-in is typically directed inwards to the fovea. However, by appropriate experimental manipulation, it is possible for the direction of filling-in to be altered (Breitmeyer & Jacob, 2012). For example, in Illustration 55, the black ring makes a strong inducer and not only fills inwards, but often outwards as well (Hamburger et al., 2006).

In this chapter we are more concerned with filling-out phenomena due to eye-motion, rather than upon fixation which reduces motion. This outward direction of induced filling-in would be a form of extrapolation (see page 57), but could be involved in the main effect of interest.

3.2   Alternative Models

3.2.1   Reason for the lack of computational models

There are no computational models for the most significant aspects of trans-saccadic perception – the fusion of information from several fixation points. The reason is that most models are conceptual while computational ones focus on other aspects of trans-saccadic perception (Hamker, Zirnsak, Ziesche, & Lappe, 2011). In this section we review the reasons for the paucity of models. In the next two sections we will review both conceptual and computational models that may be involved in supposed parts of the overall system.

3.2.1.1   Controlling motion over static scenes vs. engaging with motion over dynamic scenes

Both recurrence and continuous motion interactions are features of cognition that complexify analysis and are often ignored due to overall lack of knowledge within the field. However, if always ignored, science cannot progress. One reductionist way of ignoring the difficulties of motion, is to focus research on controlling it, perhaps in a sensorimotor fashion, while ignoring the perceptual ramifications. Causing motion is easy; dealing with the consequences, such as fusion and stability, is much more difficult. It is much simpler to assume that a planned motion to point X has been completed, then to dynamically figure out whether point X has been reached and whether motion ought to stop. Therefore, even though control of motion implies eventual motion, this controlled motion is ignored and for simplicity presumed atomic and correctly executed. Every image frame is treated as a brand new static image for which a new motion must be planned. Image frames during motion are ignored.

Tatler (2009) reviews eye guidance research. The quantitative and computational models that do exist are based exclusively on static scenes. Until natural and dynamic interactions are accounted for, current accounts may be nothing "more than descriptions of how we look at pictures" and do not relate to situated behaviours that have a tight spatial and temporal coupling while *engaged* in a cognitive task.

Our research is also not complete at this point in time. We presume fixational motion has already been planned and manage the complex ramification of continuous fusion during

continuous motion. We extend research to a new domain of inquiry – managing and fusing continuous visual sensations.

3.2.2 Computational models

There are no computational models for the most significant aspects of trans-saccadic perception – the fusion of information from several fixation points. However, there is a plethora of non-integrated computational models for the system and its supposed parts. We review a few here.

Despite the existence of incessant motion, few, if any, cognitive models of visual *perception* take eye motion into account. Models seem to be split across the vision for action or perception pathways (Goodale & Milner, 1992) despite the fact that these distinctions could not be "pure" (Cardoso-Leite & Gorea, 2010) and would continually interact or modulate each other (R. I. Wilson & Du Lac, 2011). Models with eye motion focus on sensorimotor integration at the expense of perception, e.g., (Haazebroek, Van Dantzig, & Hommel, 2011). While they are geared to controlling motion, they do not address the perceptual side-effects of motion that include issues of image stability, visual memory for parts of a scene no longer sensed, nor of mapping to an alternative coordinate system.

3.2.2.1 Integrating motion signals for predicting location, not for integrating vision

Many computational models will integrate various position or motion signals to produce a final estimate of object position. These signals could include 1) advance motor plans, 2) proprioceptive eye feedback, 3) predicted object motion from an old location, and 4)

object retinal position changes. The objects would be simple point sources or flashes at two points in time, so the integration is over a single coordinate and does not involve the trans-saccadic fusion of images. Instead, the purpose of these models is to explain peri-saccadic phenomena involving image distortions. The signals are often considered noisy. In one sense, these models are more advanced than ours in considering object motion; yet they avoid the hard problem of fusing images.

Many of these models assume the availability of a continuous internal eye-position signal, but a signal with such characteristics has not yet been found in the brain (Hamker & Ziesche, 2011). Our model currently does not make use of proprioceptive feedback and it is for future research as to how such a signal, even in a non-continuous form, would be used to minimize the error between actual and internally modelled position.

3.2.2.1.1   Optimal trans-saccadic integration

For example, (Niemeier, Crawford, & Tweed, 2003) present a Bayesian model that optimizes the use of the various signals. Their conclusion is noteworthy

> *"Our results suggest that the brain constructs its evolving picture of the world by optimally integrating each new piece of sensory or motor information."*

Could anyone disagree, especially as the priors required in defining *optimal* involve many free parameters? More specifically, Niemeier *et al.* note that their model could use either feedforward efference signals or proprioceptive feedback. This indicates that their system is not a real-time process model. Finally, a problem with Bayesian approaches in general is that they lack plausible modelling specificity and detail above and beyond the definition of optimal:

> *"The computations are complex, but they could be learned from available sensory and motor signals (they are well approximated by artificial networks of just a few hundred sigmoid neurons). We suggest that natural selection and learning shape the brain's circuits into transsaccadic integrators that make optimal use of each piece of sensorimotor information."*

It is not clear how such circuits would be distributed across the field of vision, i.e., under a multiple objects tracking paradigm. There is an assumption of object vision for this model, so it is not clear how all the pathways are integrated – it is not a biologically plausible model. Indeed, Bayesian models do not model mechanism (Jones & Love, 2011).

3.3   Emergic model

In order to explain the phenomena of this chapter, we have taken the Emergic Cognitive Model (ECM) architecture as summarized on page 5 , and made no changes.

While we call the various phenomena in this chapter as *trans-saccadic perception*, *peri-saccadic perception*, *saccadic suppression/masking*, *filling-out* and *drawing*, we have included no separate mechanisms to achieve any such behaviour – they are epiphenomenal with respect to our model.

3.4   Test 1-3: Robust stimuli

In these tests, we demonstrate the robustness of our trans-saccadic filling-out model by examining its emergic behaviour across three different stimuli.

3.4.1   Methods

The methods used are identical to those of page 36 except as indicated herein.

3.4.1.1　Subject analysis summary

This virtual agent was stochastically generated based on the development parameters detailed in the supplementary material (s43). Its stochastic retina is shown in Illustration 56. To focus characterization on the filling-out phenomenon, the field of view for the receptive field layers (208′) was made twice as large as the underlying photoreceptive field (91′) while other aspects were simplified.
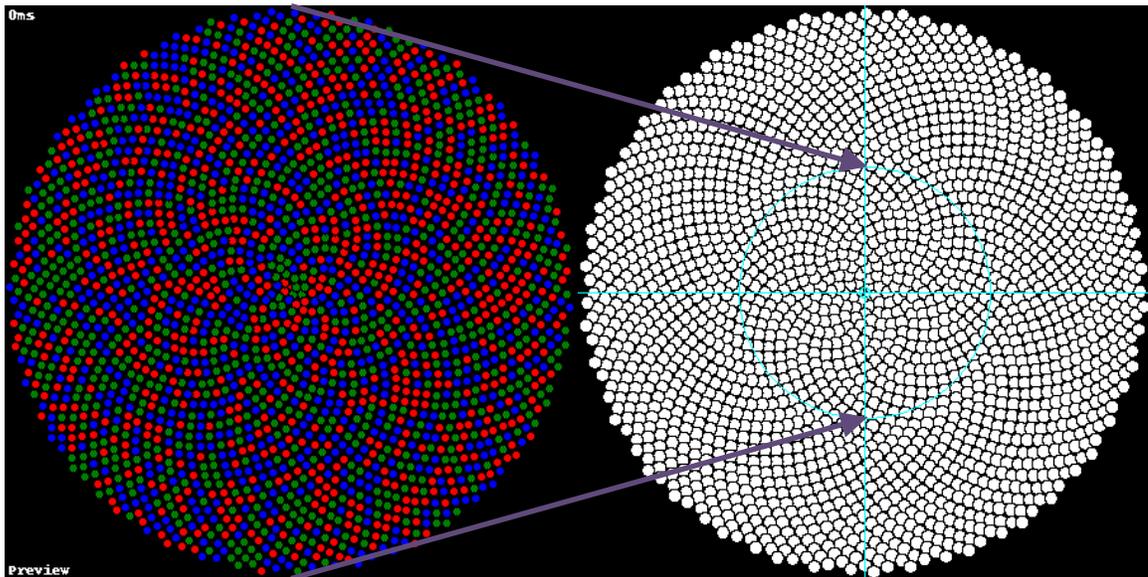
Illustration 56: Stochastic retina mapped into part of the first level RF for filling-out testing

3.4.1.2　Stimuli

Three different stimuli were used which are fully described in the supplementary material (s46, s47, s48). Each covered the entire visual field of view for the entirety of the test. For test 1, an image with four differently coloured homogeneous surfaces of various shapes – block letters spelling *Love* was used (Illustration 57). For test 2, a standard natural image – that of Lena (Picard, 1995) – was used (Illustration 58), and for test 3, a dynamic image with surface gradients – the lilac chaser (Blue & Szirko, 2005) – was used (Illustration 59).
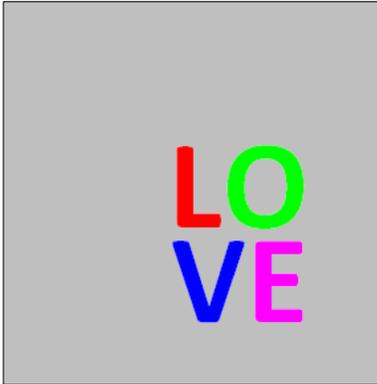
Illustration 57: Test 1 stimulus **Illustration 58: Test 2 stimulus** **Illustration 59: Test 3 stimulus with view overlay**
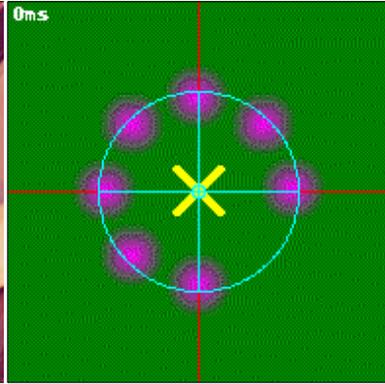
3.4.1.3   Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to move its eye in the following pattern.

- Move right for 8 ticks, target=(60′, 0′)

- Move down for 8 ticks, target=(60′, -60′)

- Move left for 8 ticks, target=(0′, -60′)

- Move up for 8 ticks, target=(0′, 0′)

The artificial brain (ECM) was monitored for the first 32 time ticks. Note that there is an eye movement planning and distribution delay before the eye actually moves as desired.

3.4.2   Results

The visual frame-by-frame response results with precise timing information are detailed in the supplementary material (s51, s62, s73). When presented as an animation (Leibovitz, 2012a), all the results for this chapter can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=86
```

In this section, we summarize and discuss various highlights. We also extract memory cohesion times.

3.4.2.1   Trans-saccadic filling-out behaviour

The dynamic behaviour of filling-out for the three tests are summarized in Illustration 60, Illustration 61 and Illustration 62. The central region, bounded by the cyan circle and axes, contains the information for the receptive field level 1 which receives upcoming information from the retinal photoreceptors. Beyond this retinal region, the receptive field can only receive information from lateral receptive fields, if the level has shifted to compensate for eye motion. Shifting starts at tick=7. While every tick is nominally 10ms, we have not yet calibrated it with biological findings. At this point in time, we are only concerned with relative dynamics.
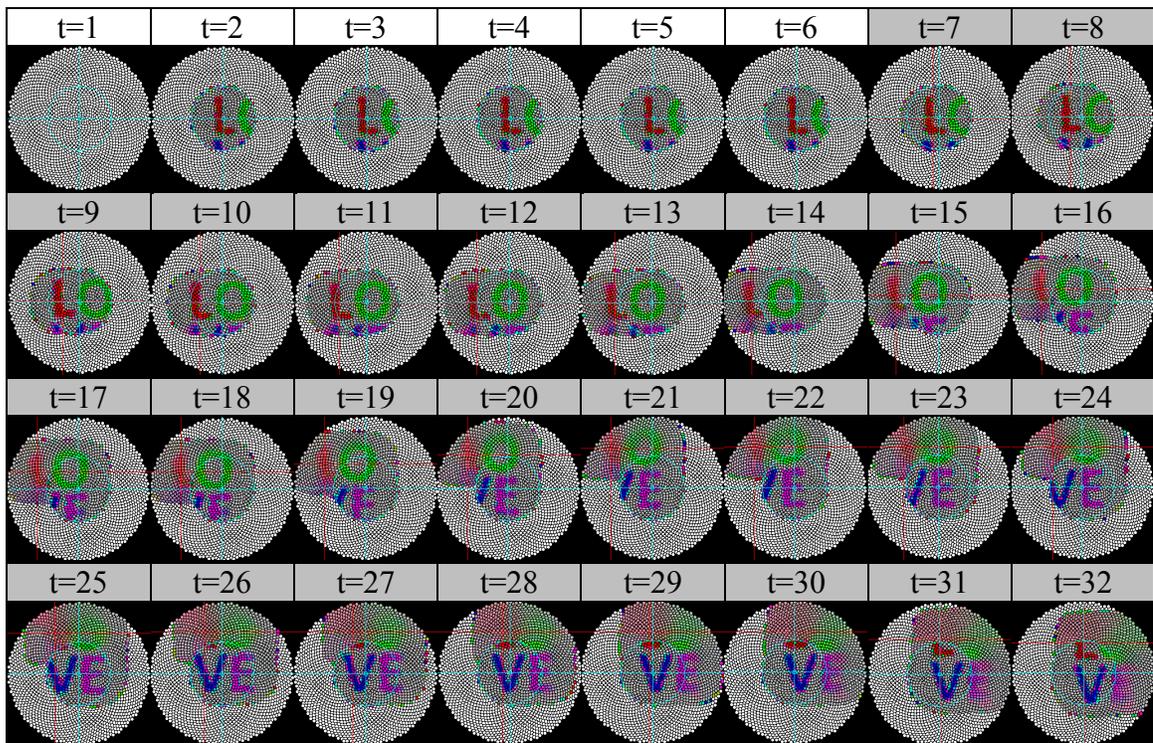
Illustration 60: Test 1 Filling-out RF level 1 with simple surfaces
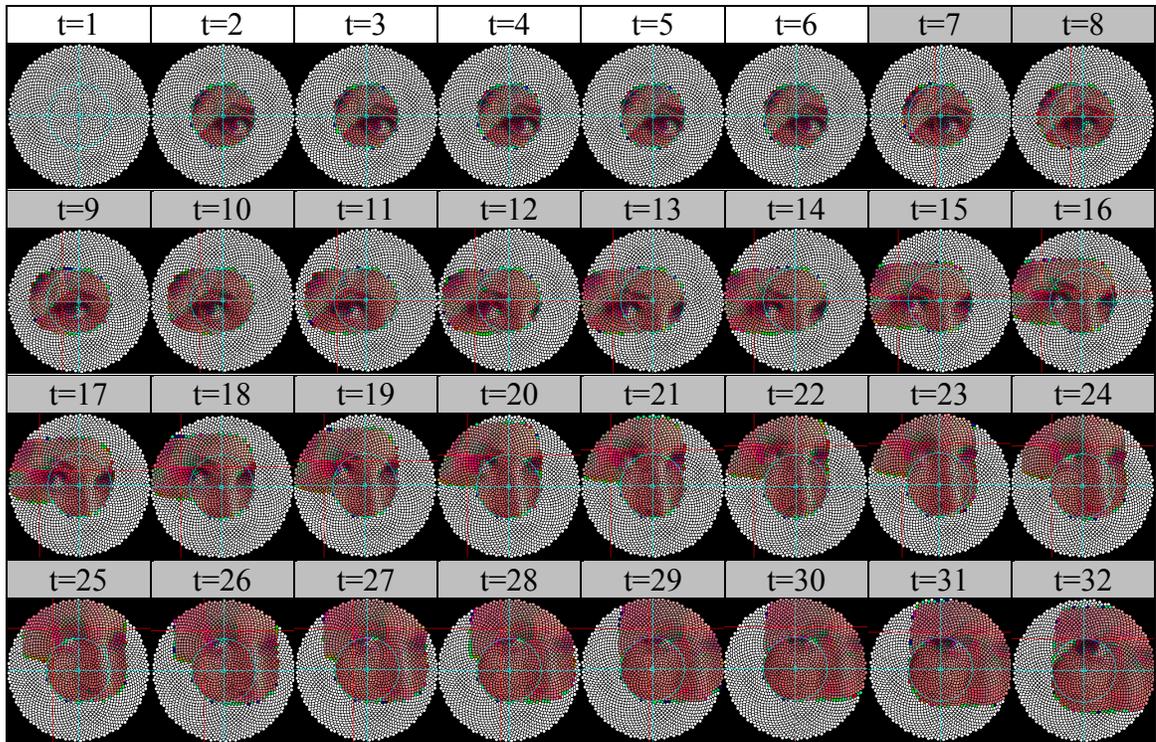
Illustration 61: Test 2 Filling-out RF level 1 with natural image
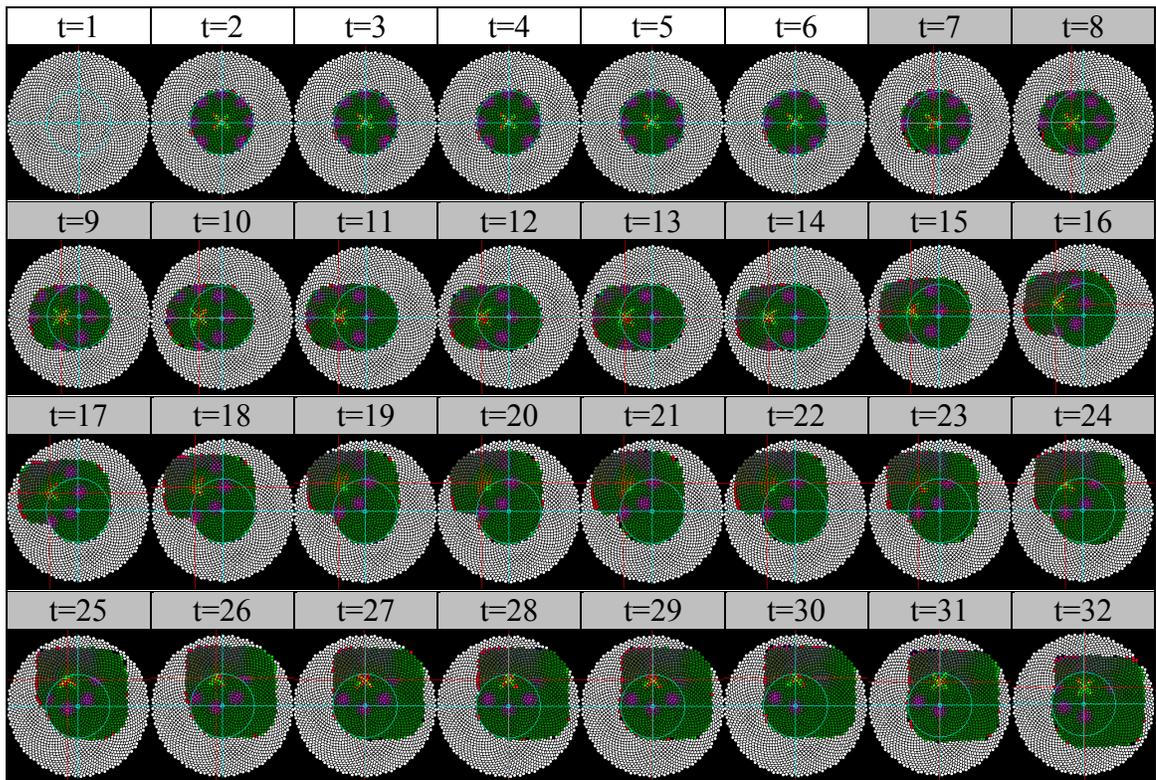
Illustration 62: Test 3 Filling-out RF level 1 with dynamic image

84

3.4.2.2 Filled-out example

Filling-out is fully demonstrated in the last frame at tick 32 as shown in Illustration 63,
Illustration 64 and Illustration 65 . While all the information is aligned seemingly on a
headcentric *canvas*, visual memory did not last long enough to leave a clear imprint on
that canvas in all cases.
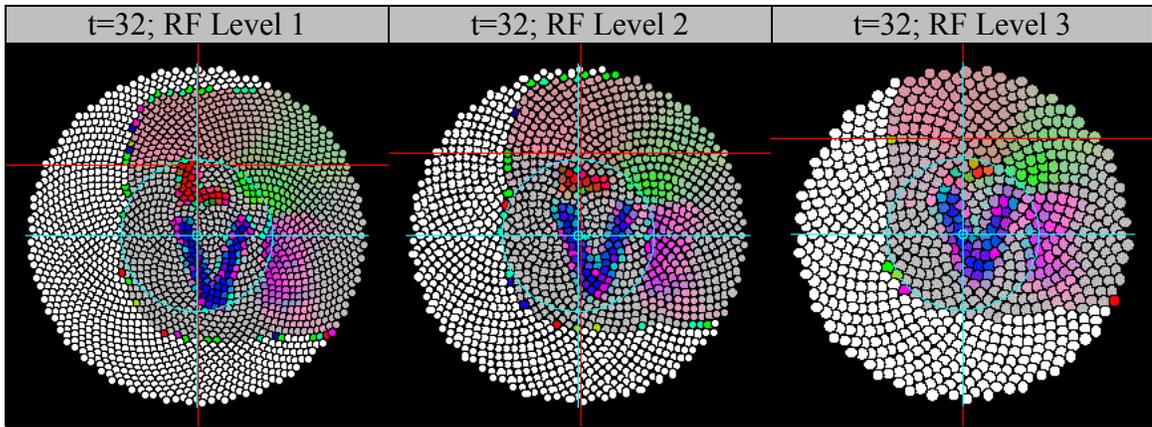
Illustration 63: Test 1 Filled-out RF levels for simple surfaces
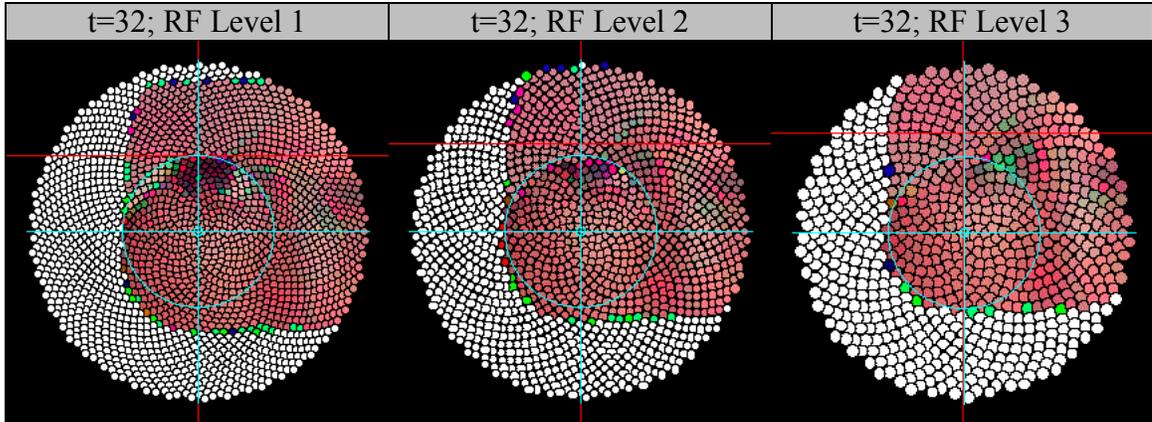
Illustration 64: Test 2 Filled-out RF levels for natural image

t=32; RF Level 1	t=32; RF Level 2	t=32; RF Level 3
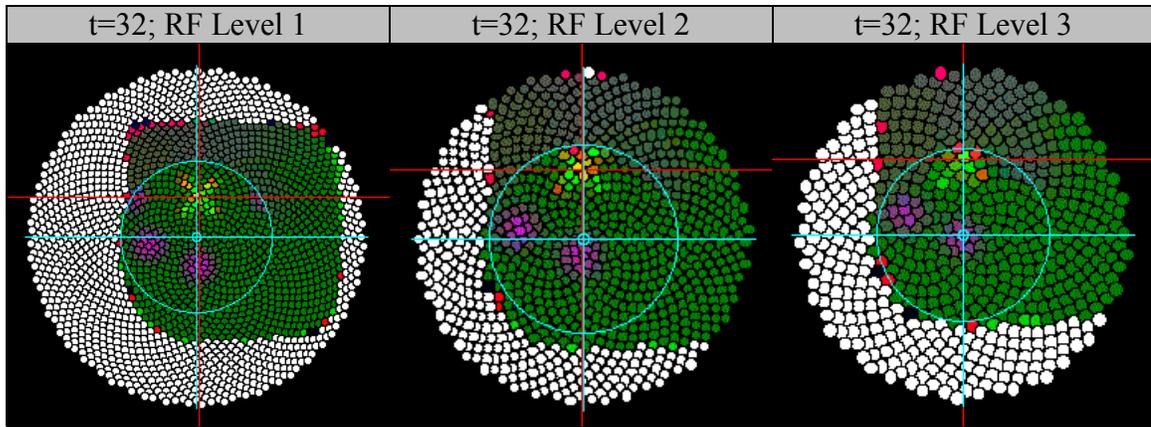
Illustration 65: Test 3 Filled-out RF levels for dynamic image

3.4.2.3 Memory performance

For test 1 (Illustration 60) at tick 9, part of letter L has moved beyond the retina. By tick 14, the entire L is completely beyond the retina and it is still cohesive (legible)

t=14	t=18
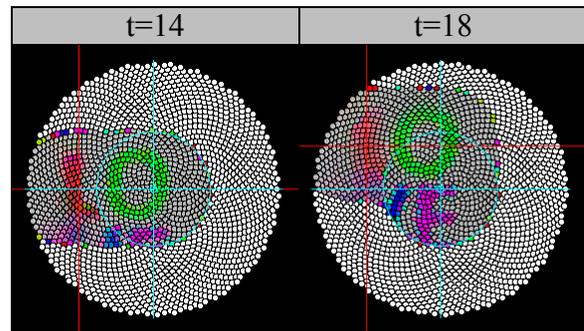
Illustration 66: Letter 'L' beyond the retina

by tick 18. These are expanded in Illustration 66. After tick 18, the 'L' dissipates into the local average scene statistics, still retaining its red surface colour properties. The red axes represent the headcentric coordinates, fixed at the same point on the 'L' throughout. The animations are shown stabilized at the (0, 0) headcentric coordinates. By tick 17, it is obvious that the eye has fully moved to the right and is now saccading down.

For test 2 (Illustration 61), part of Lena's left eye moves beyond the retina at tick 10. By tick 14, the entire eye is mostly beyond the retina and it is still cohesive (legible) by tick 23. These are expanded in

t=14	t=23
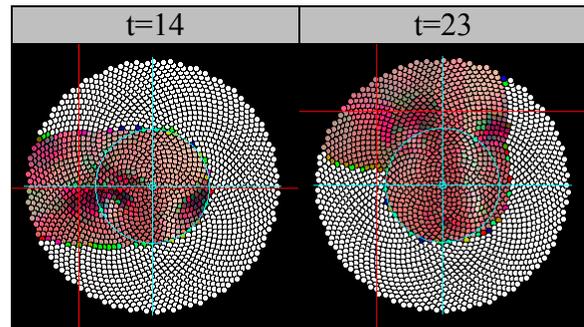
Illustration 67: Lena's eye beyond the retina

Illustration 67. After tick 23, the eye dissipates into the local average scene statistics. Filling-out is demonstrated in the last frame at tick 32 as shown in Illustration 64. All the information beyond the retina is clearly aligned seemingly on a headcentric *canvas*. Due to better memory performance, this canvas was better painted than Illustration 63 on page 85.

For test 3 (Illustration 62), the top-right lilac moves beyond the retina at tick 15. By tick 17, the entire lilac is completely beyond the retina and it is still cohesive (legible) by tick 21. These are expanded in Illustration 68. After tick 21, the lilac
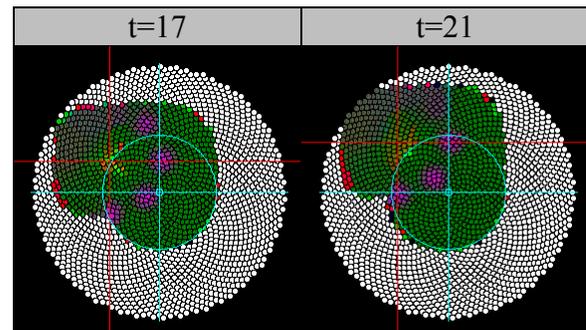
Illustration 68: Lilac beyond the retina

dissipates into the local average scene statistics, still maintaining the lilac colour. In this particular test, eye motion was lessened and it becomes clearer that the retina "paints" the "canvas" forming the outer regions. For example look at the final image (t=32) of Illustration 62.

3.4.3   Discussion

3.4.3.1   Loss of detail

For test 1, the filled-out information has a short memory (9 ticks) and does not look like a stitched together panorama by tick 32. It has lost a great amount of detail, although some statistics such as colour may remain longer. While not of satisfactory appearance, this is an important point in its biological plausibility. Change detection studies originally ruled out the concept of a detailed visual buffer because many changes to details, especially

when coinciding with saccades, blinks or other interruptions, could not be detected –

scenes can either be coherent or detailed, but not both (Rensink, 2002). So suggestions of

gist-like or proto-objects have been made. However, change blindness need not imply

that sparse representations are mandatory, as detailed information can last up to 0.5s

(Simons & Rensink, 2005) and other possibilities could also explain the inability to detect

changes (Tatler & Land, 2011). Indeed, latter studies show that visual memory can

accumulate and be detailed (Hollingworth, 2006), and that the change blindness results

may be due to insufficient encoding or attentive times (Brady, Konkle, Oliva, & Alvarez,

2009).

For test 2, the filled-out information for the Lena stimulus has a longer memory (13 ticks)

than for the Love stimulus. Moreover, the image always has a properly stitched together

trans-saccadic form. This is likely due to less contrast among visual features within Lena

that minimize interference, but this has not been quantitatively ascertained.

For test 3, the filled-out information for the Lilac stimulus has a short memory (6 ticks)

due to dissipation that merges one lilac disc into another. Nevertheless, their spatial

location is maintained and the image always has a properly stitched together trans-

saccadic form.

3.4.3.2   Diffusion vs. cut & paste shifting

While we call this account filling-out, it is not a form of filling-in where information

diffuses outwards. Even by tick 32, the left third and bottom third of the receptive field

level is still empty – nothing has diffused into it. That's because there is no diffusion in

our account. We shift to compensate for eye motion, and the eye has not moved sufficiently to cause a shift downwards or to the left.

A better analogy to shifting is to imagine that the central retinal connected region is a paintbrush, but that rather than moving the brush, the canvas is moved around instead – the canvas is always moved in the opposite direction as the eye. The shifting effectively amounts to a form of *cut & paste* (Van Roosmalen et al., 1999) rather than any form of diffusion.

3.4.3.3   Flowcentric Memory

The Lena stimulus is ideal to demonstrate the *flowcentric* nature of our information processing approach. In Illustration 69, one can clearly see that Lena's eyes are centered within the 1st Level RFs
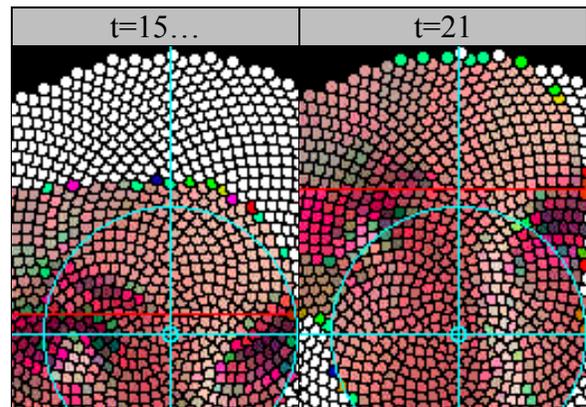
Illustration 69: Flowcentric memory

corresponding to the photoreceptors of the retina (cyan circle) at tick 15, while they are the just above the retinal region at tick 21. This will also occur at the 2nd and 3rd RF levels. The point is that the full memory of Lena is not associated with any neurons either in the retinotopic or spatiotopic reference frames. While her eyes are over the photoreceptor region, they form *direct* retinocentric perception which is highly transitory and need not be considered as memory. It is when they move beyond the view, that an *indirect* memory of Lena must exist somewhere, using the terminology of direct perception (Gibson, 1950).

The question is where does the indirect memory of Lena reside, and is memory even an applicable construct? To answer the first question, the memory of Lena resides in the continuous flow of information. Lena is perceptually stable within the flow as indicated by the red axes. One can make an analogy to the flow of water in a river. One can drop some coloured oils onto the surface of the water, painting a simple picture if you will, and this picture will maintain cohesion for some time. It does not really matter if the river turns left or right, or if it even narrows and moves more quickly. Within the flow of water, the picture maintains its relative characteristics. Of course, the picture will continuously change in an unexpected manner, and in this sense, perhaps it should not be called a memory. This would be unlike *neurocentric* memories where the image might become fainter, but will not lose its spatial precision.

The purpose of ECM's *maintain information* coherence function is to ensure that information processing within the brain/mind maintains flow coherence throughout regardless of the current set of neurons keeping the flow moving. We think of neurons as paddle wheels keeping the water *computationally* flowing – it does not really matter where these water wheels are placed. This analogy is similar to the internet. It is the relative flow of packets that is important regardless of which routers they pass through in order to get to their destination. It does not matter too much how many routers there are (beyond delay times), where these routers are, and whether each packet follows the same route. This does not mean that one shouldn't analyse routers, but that an analysis of the data flow may be equally important.

Conceptually, at the boundary of sensors and effectors, the internal flow of information is used to engage with a flow of external interactions. In this sense the flow is always spatiotopic.

3.4.3.4 Coordinate systems

The filling-in behaviour is visually summarized (Illustration 60) in retinocentric coordinates, but detailed in the supplementary and animated (http://emergic.upwize.com/?page_id=86) in headcentric coordinates that appear to show image stability.

3.5 Test 4: Edge under varying RF sizes

Image stability often implies a mapping from one coordinate system to a motion invariant one. In the previous experiments all the RFs were of the same size simplifying such transformations. In this test, we characterize the effect of more complex retinal topologies. We hypothesize that ECM is invariant to physical topology.

3.5.1 Methods

The methods used are identical to Test 1 except as detailed herein.

3.5.1.1 Subject analysis summary

A virtual agent with similar characteristics to those of Test 1 was used except that the smallest and largest cone size sizes varied by a factor of four. The detailed parameters are available in the supplementary material (s45) and the results shown in Illustration 70.
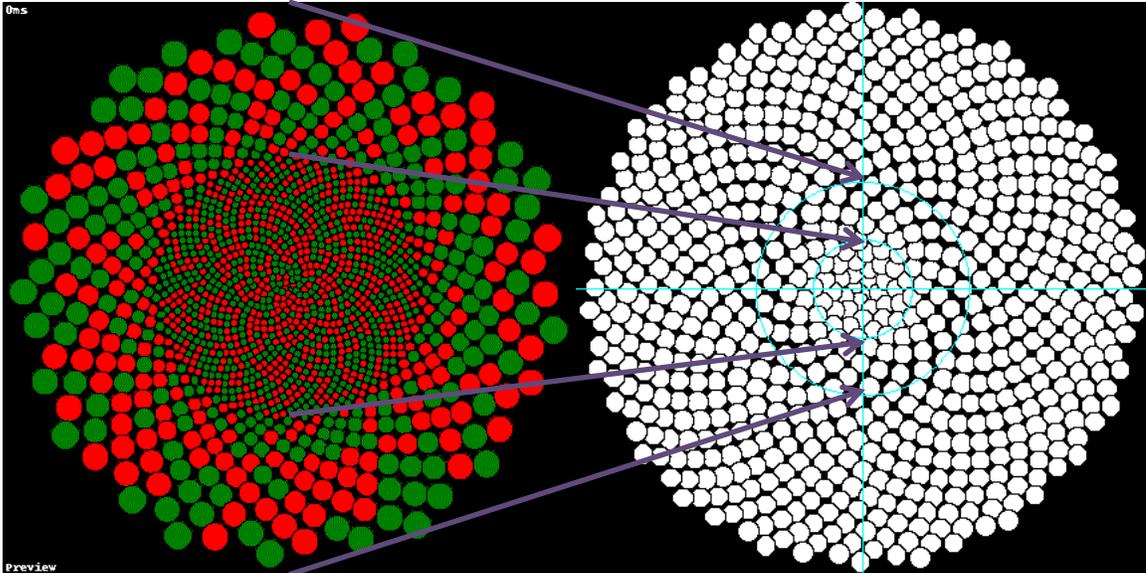
Illustration 70: Stochastic retina with varying RF sizes mapped into 1st level of RFs

3.5.1.2 Stimuli

The cones of the ECM retina develop in a circular sunflower pattern. Converting from such a topology to a rectilinear one (and vice versa) can introduce artefacts, for example, a straight edge may not appear straight after transformation. We used a single vertical green/red edge to test the accuracy of any transformations. The stimulus
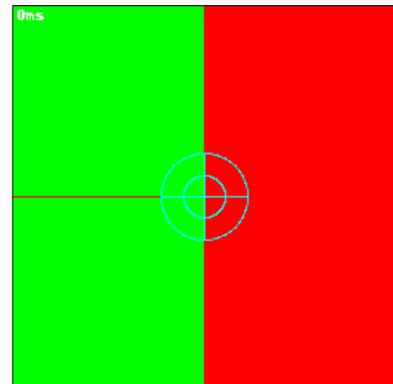
Illustration 71: Test 4 stimulus with retinal overlay

covered the entire visual field of view for the entirety of the test. The stimulus is fully described in the supplementary material (s49), and appears in Illustration 71.

3.5.1.3 Procedure

The procedure is similar to that of Test 1. Movement was in four 60′ increments in the directions up, right, down, left.

3.5.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s84). In this section, we summarize and discuss various highlights. We also extract memory cohesion times.

In Illustration 72, the straight edge first makes its way beyond the retina at tick 7, and maintains its straightness till tick 22 (and does not lose cohesion by tick 33). These are shown in expanded form in Illustration 73.
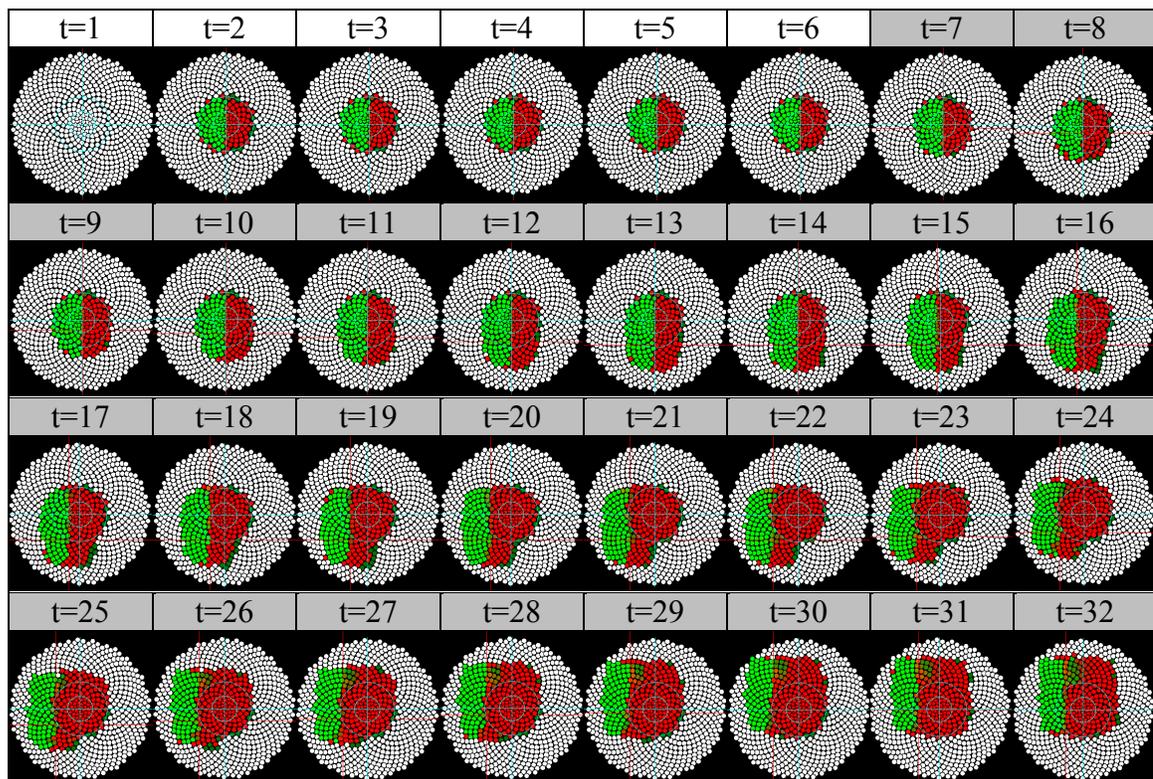
Illustration 72: Filling-out RF level 1 with a straight edge

The red axes represent the headcentric coordinate which happens to fall on this straight edge. Note that the edge cannot be perfectly straight due to underlying patches.
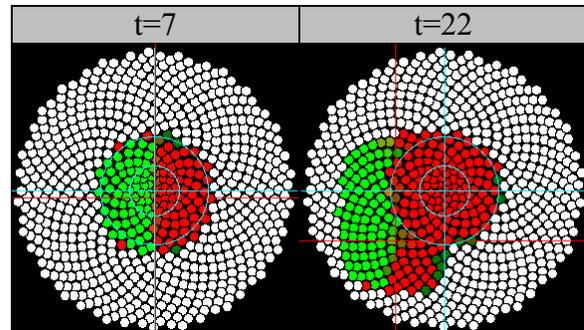
Illustration 73: Edge beyond retina at RF level 1

3.5.3 Discussion

The straight edge appears straight at all times confirming the hypothesis that the underlying topology makes little difference to the coordinate transformation system underlying ECM. ECM seems invariant to topology. The reasons are as follows. Firstly, the GUI shows the RFs in linear coordinates. In particular, the size of the RFs shown is their functional size in terms of their field of view. It would not be in their physical size which is fairly constant outside the foveola. Thus no log-polar transformation from physical brain to Cartesian space is required, e.g., as performed in (Fazl, Grossberg, & Mingolla, 2009). Secondly, every receptive field knows its coordinates and only accepts values (with statistical manipulation) tagged with those coordinates. An RF will broadcast its value laterally to a local population. As long as eye movement is not too great, the appropriate RFs will receive the value regardless of whether they are larger or smaller. ECM does not perform coordinate transformation across reference frames. RFs always maintain their retinocentric coordinates shifted to compensate for eye motion, while values always maintain their headcentric coordinates that are never changed. The value may have to flow laterally to compensate for eye motion, but its coordinates are not changed.

Note that the visual system is not fully shift invariant (Schwartz, 1983), so some distortions are to be expected.

3.6 Test 5: Image masking under fast motion

The concept of active, extra-retinal image suppression fails when the eye can capture moving stimuli that matches eye motion (García-Pérez & Peli, 2001; Kelly, 1990). It is the passive, retinal information that could account for masking instead. While ECM does not yet handle image motion, nor image blur, we simulate the effect in this test by moving the eye at varying speeds. We hypothesize that information will be effectively masked when the eye moves too fast, i.e., the masking phenomenon will emerge in ECM as required.

3.6.1 Methods

The methods used are identical to Test 1 except as detailed herein.

3.6.1.1 Subject analysis summary

The agent is the same as in Test 1, but the following details are relevant. The size of each cone is 3.68′, and along with various fan-out factors, this will affect the speed at which information is effectively masked.

3.6.1.2 Stimuli

A wide stimulus was chosen to allow for large horizontal eye movements (s49). It is shown within Illustration 74.
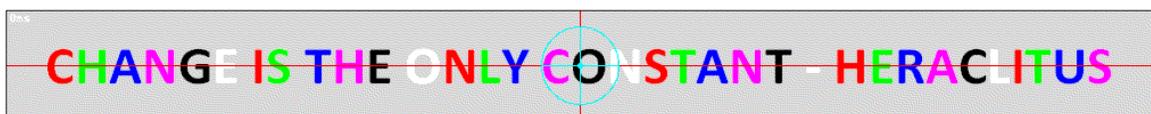
Illustration 74: Test 5 stimulus with retinal overlay

3.6.1.3   Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to accelerate its eye in the following pattern.

- Accelerate right by $+1.5'/\text{tick}^2$ for 16 ticks

- Accelerate right by $-1.5'/\text{tick}^2$ for 8 ticks

The artificial brain was monitored for the first 24 time ticks. Note that it there is an eye movement planning and distribution delay before the eye actually moves as desired.

3.6.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s95). In this section, we summarize and discuss various highlights. Unlike the other experiments, we do not extract memory cohesion times as we are trying to mask memory.

At tick=17, information at RF level 1 exceeds its lateral fan-out and starts to be effectively masked. This can only be seen in the portion beyond the photoreceptor view returning to a null state. By tick 20, information is completely masked. Thus, it takes 3 ticks for masking to complete as shown in Illustration 75. The speed is shown in cones/tick diameter so it can be related to fan-out.

t=16	t=17	t=18	t=19	t=20
Speed=4.08	Speed=4.48	Speed=4.89	Speed=5.30	Speed=5.71
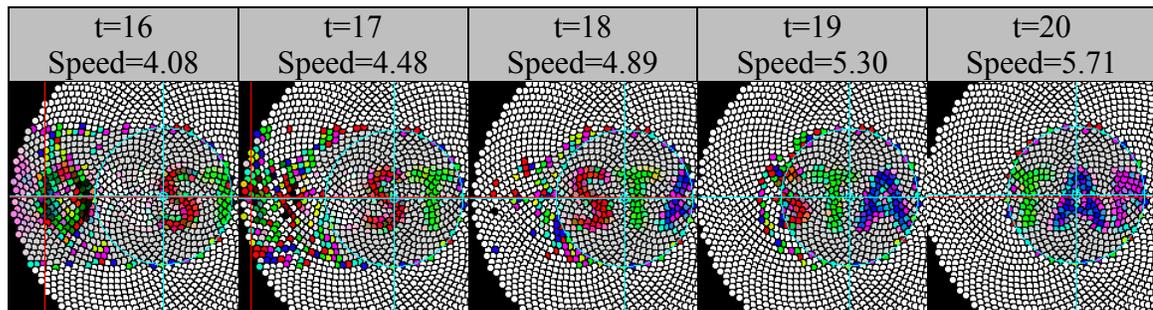
Illustration 75: Masking at RF level 1

For level 2, the masking starts at tick=18 and completes by tick=21. It takes 3 ticks for masking to complete as shown in Illustration 76.

t=17 Speed=4.08	t=18 Speed=4.48	t=19 Speed=4.89	t=20 Speed=5.30	t=21 Speed=5.71
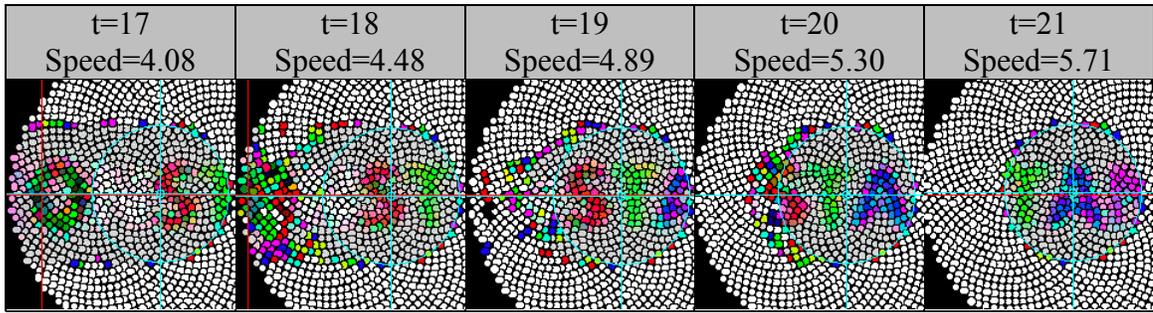
Illustration 76: Masking at RF level 2

For level 3, the masking starts at tick=19 and completes by tick=22. It takes 3 ticks for masking to complete as shown in Illustration 77.

t=18 Speed=4.08	t=19 Speed=4.48	t=20 Speed=4.89	t=21 Speed=5.30	t=22 Speed=5.71
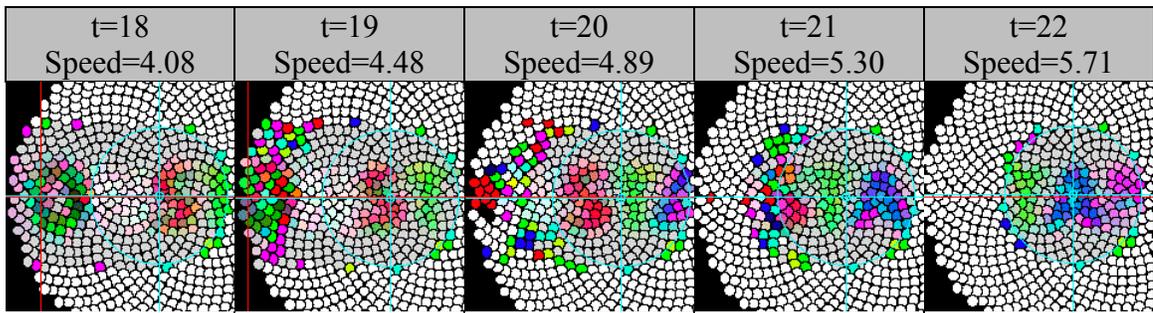
Illustration 77: Masking at RF level 3

3.6.3 Discussion

To maintain biological plausibility, RFs are only connected to a local neighbourhood as specified by various fan-out factors. Information flows across these links to compensate for eye motion and ensure information coherence. If the information flows too fast, it is effectively masked as there are no receptive fields to capture the flow at the new destination. Thus, this behaviour mimics the retinal, passive masking phenomenon, and supports the hypothesis that such phenomenon can emerge in ECM.

3.7   Test 6: Drawing figures via planned eye movements under paralysis

Subjects under eye muscle paralysis dissociate the updating of trans-saccadic memory based on planned eye motion, with actual eye motion. When looking at an unmovable spot of light and by trying to saccade to the left, they would perceive that point source as moving to the left even if their eyes did not move (Whitham et al., 2011). Three of the four subjects could control these illusive movements well enough to draw

Illustration 78: Drawing figures via planned eye movements under paralysis. From (Whitham et al., 2011), Fig. 3. By attempted eye movement, a point source of light could be controlled to draw simple figures (A-D), and (E) even sign a subject's name "Dy[lan]". The trajectories are recreations and were not perceived themselves.

simple figures and sign the first two letters of their name – see Illustration 78. In this test we recreate these findings of filling-out from a point source of light.

3.7.1   Methods

The methods used are identical to Test 1 except as detailed herein.

3.7.1.1   Apparatus

The software was engineered to inject a stimulus into the 10 innermost photoreceptors. This was eye-position invariant.

3.7.1.2   Stimuli

The agents were disconnected from the view so that none of the photoreceptors received any stimuli, except for the 10 innermost photoreceptors that are arranged in a circular pattern. They were sent a value of 127/255 which translates to a 50% grey to mimic a

point source of light. Grey was chosen because the GUI indicates the absence of colour information, e.g., NULL, in white.

3.7.1.3 Procedure

As in Test 1, but an extra 5 ticks were provided to enable the completion of the square figure. That is because motor planning exacts a 6 tick delay.

3.7.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s104). In this section, we summarize and discuss various highlights. We also extract memory cohesion times.

3.7.2.1 Filling-out/drawing behaviour

The dynamic behaviour of filling-out, or *drawing* in this case, is summarized in Illustration 79. Compare to Illustration 78.
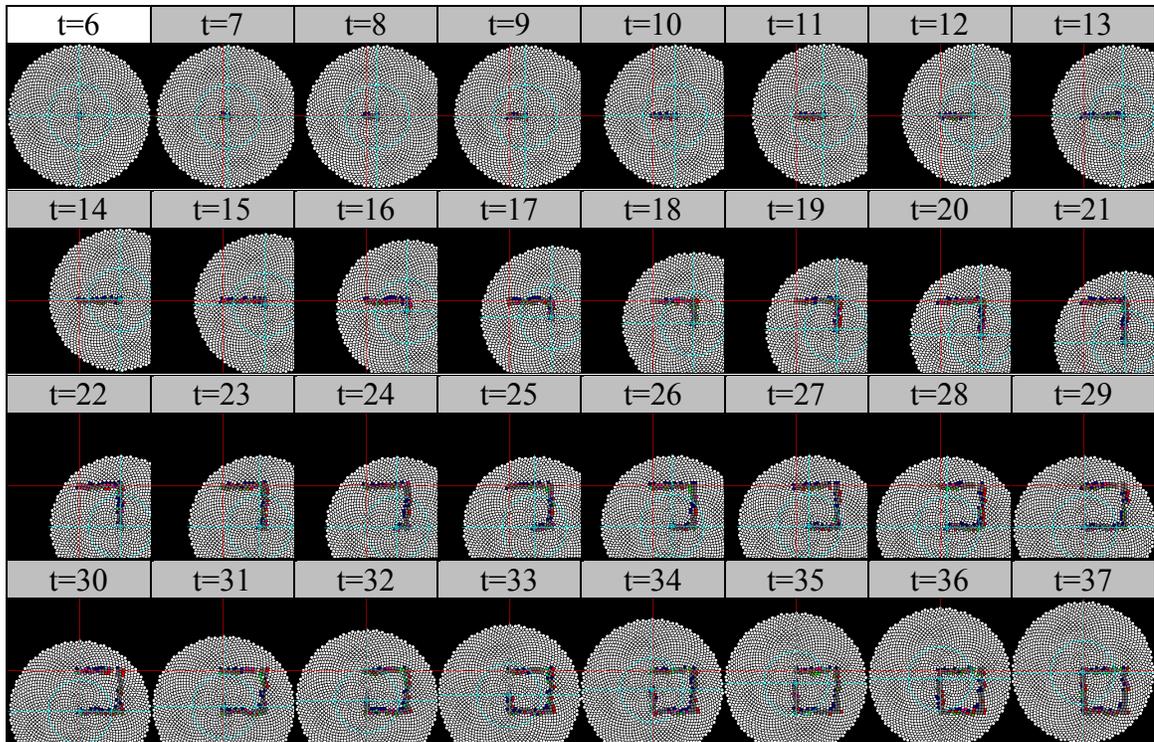
Illustration 79: Filling-out RF level 1 via drawing shown in retinocentric coordinates

3.7.2.2   Filled-out and fully drawn example

Filling-out is demonstrated in the last frame at tick 37 as shown in Illustration 80.
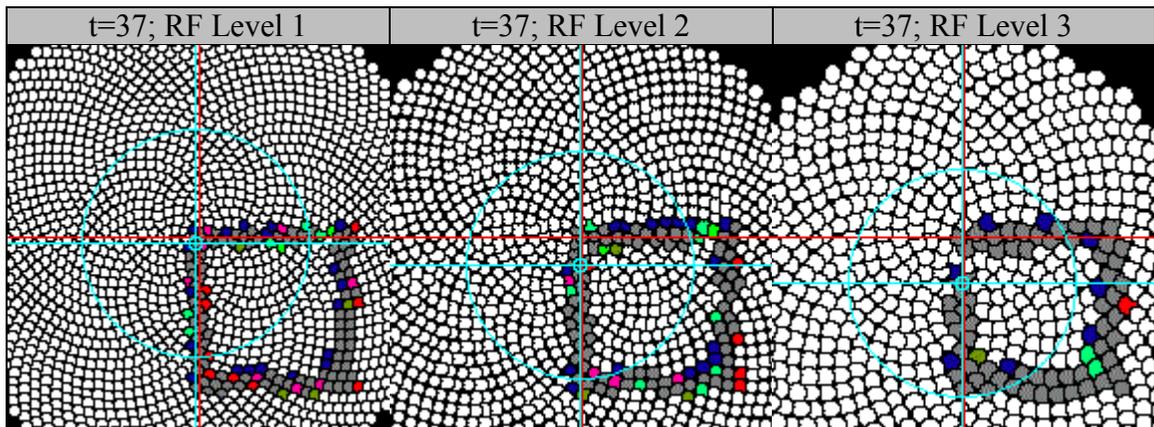
Illustration 80: Filled-out RF levels for drawn figure

3.7.2.3   Memory performance

In the previous experiments, the start of memory was considered when the corresponding part of the image moved beyond the reach of the retina. However, in this case, the agent was disconnected

Illustration 81: First and last memories for the central spot at RF level 1

from the view and could only sense the fixed central spot. Any deviation from center represents memory as there is never a square stimulus, only a point sourve. Therefore, at tick 7 the spot has first moved beyond the center and by tick 36 it is still cohesive and did not dissipate – there was no interference from any other spatial locations. This is shown in Illustration 81.

3.8   General results and discussion

In this chapter, we have seen the emergence of the filling-out phenomena despite the lack of an explicit filling-out mechanism in what appears to be a headcentric map beyond the reach of photoreceptor signalling. Thus a retinocentric to headcentric mapping system emerges despite the lack of an explicit mechanism to do so. It appears robust against retinal topologies. Image features can be retained within this headcentric map for some time before dissipating, which implies that a visual short term memory emerges despite the lack of an explicit memory system. The fading of memory could be characterized as a diffusion process towards the local image statistics, although simple edges seem to last much longer. This characterization is also emergic. The animated images clearly show that image stability is maintained (centered on the red axes), again without a dedicated mechanism.

A cognitive computational model of peri- and trans-saccadic sensory integration has been demonstrated where retinal information can be used to fill-out non-retinal locations. This is due to the lateral shifting of information to compensate for incessant eye motion – the sole purpose of the *maintain information coherence* function. Note that spatial coherence applies to the dynamic flow of information, so we call this a *flowcentric* or *infocentric* paradigm. Because image stability and space constancy are considered to be properties of conscious perception, and because both properties are found within the flow, we hypothesize that the locus of consciousness is to be found in the flow.

3.9   Derivative narrow slit viewing model

The trans-saccadic model fuses many small spaces over time to produce a larger spatial representation. This can be adapted to anorthoscopic viewing whereby the world is viewed from a moving narrow slit, and a larger representation is built up. Such a derivative model is described in Appendix A: An emergic model of filling-in for anorthoscopic perception on page 201.

4 Chapter: An emergic model of filling-in after brief stimuli

Typically, recurrence leads to explicit or emergent forms of memory (Maex & Steuber, 2009), but even purely bottom-up visual models will have memory effects if they have multi-scale temporal filters (e.g., Baraas, Kulikowski, & Muldoon, 2010). Memory is unavoidable in any physical process that takes real time to complete. Within the retina, transient and sustained bipolar cells represent high and low temporal frequency subdivisions (Valderrama-gonzalez, McGinnity, Maguire, & Wu, 2009) so at the very start of visual processing, some paths already retain information longer than others – the slower path forming a memory for the faster one. Even in a single processing path, a previous stage could form a memory for a subsequent stage (Maex & Steuber, 2009).

However, in dynamic vision, memory is problematic as it can lead to motion smear (Tong, Stevenson, & Bedell, 2008). Indeed, this is one of several reasons that the folk-
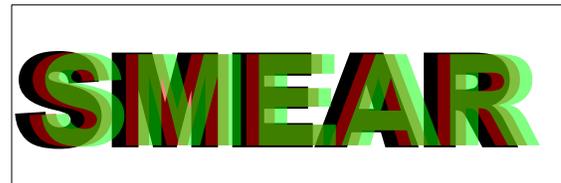
Illustration 82: Smear due to piling up of shifted black, red and green images

scientific concept of the *persistence of vision* is considered a myth (J. Anderson & Anderson, 1993) – new images will simply pile on top of old images leading to overall blurred perception as shown in Illustration 82. In this chapter, we will investigate the triumvirate of visual memory, motion and smear in relation to traditional *neurocentric* approaches to computational modelling and the *flowcentric* approach of the Emergic Cognitive Model (ECM) architecture (p323).

The tests in this chapter focus on memory after a briefly presented or "flashed" stimulus has been removed and while the eye still undergoes moderate motion. In such a scenario there are no new stimuli to pile-up and smear the old "flash memory" – there is no temporal interference. Thus we can measure the duration of memory and the timecourse of its dissipation with minimal interference, Nevertheless, it is the moderate motion that distinguishes models that merely *hold* visual information within memory (and may have trouble with smear) and those models that *maintain* memory, via *extra-retinal* updating for example, such that they effectively represent an internal model of the world – we call these *memory in motion*. These are less likely to have trouble with smear (Bedell, Tong, & Aydin, 2010). Regardless, we will compare the results in this chapter to other chapters where the identical stimulus under motion remains, and smear could be a problem – an interference between stimulus and memory. Ideally, we want a model that exhibits memory but not smear. We believe that ECM is the first such model.

Due to lack of equivalent alternative models, we will compare ECM to a model that has an explicit mechanisms devoted to handling image smear but still cannot handle motion.

4.1 Introduction

Some terminology needs to be clarified at the outset.

4.1.1 Terminology

Blur and smear are often conflated or used interchangeably (e.g., Tong et al., 2008), so in this chapter we will distinguish them by mechanism and their avoidability – after all, the main thrust of this chapter is

Illustration 83: Image blur via camera.
Image by (E01, 2008).

the avoidance of smear by appropriate use of memory and compensation for motion (prediction). Talk of such phenomenology invariable makes use of the eye-as-a-camera metaphor, and invariably exemplifies such phenomenology via computer generated images, or pictures taken from a camera as in Illustration 82 or Illustration 83. But the camera metaphor is fraught with danger (Haber, 1983; Holcombe, 2009), especially in relation to enactive tradeoffs, and concepts such as frame rate or film-speed. Similar problems occur with discrete sampling within cognitive processes (Turvey, 1977). Heading these warnings, we now define our use of these terms within this chapter.

4.1.1.1 Motion blur is physical and unavoidable

It is hard to pin down concepts such as blur (Watson & Ahumada, 2011). For the purposes of this chapter, motion *blur* is distinguished by its unavoidable physical cause relating to the limitations of the sensory apparatus, while *smear* may be avoidable and caused by information processing.

Some minimum blur does occur due to optical limitations and due to the minimum size of the photoreceptors, and it would be perceived within static images. However with fixational eye movements (such as jitter) over a static scene and along with spatiotemporal integration, such static blur could in principle be reduced (e.g., Yi, Jiang,

& Mallen, 2011). It is not clear whether human cognition takes advantage of this de-jitter capability. Some research suggests that it does not (Packer & Williams, 1992) while other research suggests that it does (Rucci, Iovin, Poletti, & Santini, 2007). Additional research suggest that jitter is compensated via extraretinal means (Stevenson, Raghunandan, Frazier, Poonja, & Roorda, 2004). In any case, the important point is that fixational eye movements did not worsen performance in cognitive acuity tasks. Motion need not lead to increased blur further suggesting that smear may be a modeling artefact.

While minimum blur over static scenes is typically less than one photoreceptor in width and defines visual acuity (Westheimer, 2009), in this chapter we are only concerned with blur and smear due to motion, and larger than one photoreceptor. The sensory apparatus effectively integrates information over a small timeframe and during this time, an object may have moved across several photoreceptors. Thus the object will appear blurred and spread across space. Motion blur is a distortion across space due to processing over time. Effectively, integration occurs within a memory, although the analog nature and access of photoreceptor memories differs greatly from those in cameras.

The time constants of the memory integration activities are therefore extremely important. Hecht & Shlaer (1936) found that at high illumination intensities, the critical flicker fusion frequency reached 60 Hz. This implies that the slowest time constant within the entire visual processing hierarchy, from sensory cone to conscious awareness cannot be larger than 17ms (assuming linearity). Physical blur need not be greater than the distance it takes for an object to move within 17ms, obviously dependent on the relative

speed of the object with the eye. Any blur greater than this we will attribute to smear and information processing – processing that is often nonlinear.

Even though blur is a physical limitation, there are physical mechanisms to prevent blur from entering the visual processing system in the first place, but these always involve tradeoffs. For example, if a camera was panned to match the speed of the bus in Illustration 83, then the bus would appear sharp, but everything else would be blurred. The panning would be equivalent to human smooth pursuit. But human vision has an advantage. It can devote attention just to the bus, and the background blur might never be noticed. Indeed the background or other stationary objects are not perceived as moving (Hafed & Krauzlis, 2010).

The shifting of attention may not suffice to explain the loss of smear due to eye motion. It is suspected that this occurs because extra retinal information is used to compensate for such motion (Bedell, Chung, & Patel, 2004). That is what the ECM architecture does, and that will form the paradigm for the tests in this chapter. In general, humans do not consciously perceive blur (Burr & Morgan, 1997), unless motion is very fast or unexpected, because of the enactive nature of vision.

There are also post-processing techniques that can reduce both blur and smear, but most of these are for future analysis. A few of these will be mentioned only because they are positioned in our area of concern.

4.1.1.2   Motion smear is avoidable information processing

While a feedforward linear analysis of the entire hierarchical visual system suggests a maximum memory integration time of 17ms, temporal sensitivities can be as low as 2ms (Sweet, 1953; Westheimer & McKee, 1977), gaps as short as 5ms can be detected (Georgeson & Georgeson, 1985; G. Smith, Howell, & Stanley, 1982), offset latencies are shorter than onset (Bair, Cavanaugh, Smith, & Movshon, 2002; Tadin, Lappin, Blake, & Glasser, 2010), and it is well known that there are a variety of persistency effects ranging from 50-1000ms (Adelson, 1983; Graziano & Sigman, 2008). Obviously the cognitive system is not linear, and both fast and slow mechanisms must collaborate (Holcombe, 2009). Ideally, the interactions are specified explicitly, e.g., via a computational model. There are likely many multi-scale paths, lateral interactions and vertical interactions. Currently, there are major disagreements on mechanism (Blake & Lee, 2005).

The problem with the myth of the persistence of vision is that it conceptually placed a slowly decaying memory store with a time constant greater than 17ms on a single feedforward path between sensation and perception. This would have led to an unacceptable amount of image blur. A similar problem occurs with computational models containing any form of memory, either implicitly or explicitly.

Thus we define smear as a blur effect caused by the information processing of computational models above and beyond their sensory apparatus. It must be avoidable as such smear does not occur phenomenologically (Burr, 1980). In this chapter we will demonstrate how *neurocentric* models invariably cause smear, and discuss why some of

their workarounds are unworkable. We also demonstrate how *flowcentric* systems do not have such problems, and that perhaps *memory* is a confusing term to describe their memory-like emergent behaviours in a world used to the storehouse metaphor of memory. Alan Turing's 1936 notion of *effective memory* (Dresner, 2012) is in line with ours.

4.1.1.3   Image dissipation is information loss not due to motion

Both blur and smear are forms of information loss due to motion. We define *dissipation* as any loss of fidelity for reasons other than motion. Ordinarily, visual memories are considered to decay
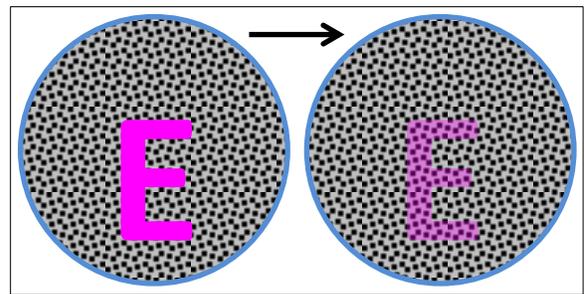
Illustration 84: Actual memory fading away

and fade over time (Zhang & Luck, 2009) and it may be hard to distinguish a faded image from a blurred one. Such memory decay would occur with or without motion and is part of *dissipation*. An example of "pure" fading is shown in Illustration 84 and does not demonstrate any blurring effects. However, models with diffusion mechanisms would be harder to distinguish. Some models may be distributed. The partial reports for iconic memories degraded when a position or colour cue was delayed, but not for a semantic category cue (letter or digit). Thus different aspects of the icon may be dissipate (Coltheart, 1980).

4.1.2   Memory in motion

The tests in this chapter will characterize the duration of flash memories, and their compensation for motion, often considered from a predictive point of view. In this

section, we review the duration of various phenomenal memories, as well as aspects of memory in motion that include predictive capabilities.

4.1.2.1 Genesis of memory

When models of visual processing have any form of recurrence or temporal delays then they may exhibit behaviour associated with having memory (Maex & Steuber, 2009). In *neurocentric* models, such memory is explicit and contained within the neuron. In dynamic models, such memory may also be explicit and is contained within the relative phase dynamics (Assisi, Stopfer, & Bazhenov, 2011), often still associated with one or two neurons. For example, the first digital computer memory was the flip-flop – a form of bistable mulitivibrator, and cognitively plausible versions that make use of spiking neurons have been realized (Colliaux, Molter, & Yamaguchi, 2009; Kleinfeld, Raccuia-Behling, & Chiel, 1990).

Explicit forms of memory realize the archival or storehouse metaphor – a memory trace is stored at one point in time to be perfectly retrieved at some arbitrary later time (assuming it hasn't decayed) (M. R. Bennett & Hacker, 2001; Brockmeier, 2010; Koriat & Goldsmith, 1996). This storage metaphor is not taken literally by all (Dudai, 2009; Echterhoff, 2011; D. C. Palmer & Donahoe, 1992), but it is inherent in the static snapshot view of visual processing. It may work well for long term visual memories (Brady, Konkle, & Alvarez, 2011). However for short term or working purposes, it may cause artefacts such as smear when the snapshot metaphor fails under continuous motion. Instead, this chapter will demonstrate an implicit *flowcentric* or *infocentric* model of "memory". In this form, information is continuously received, transformed and

transmitted by neurons – there is no delayed reading of an old memory. Indeed, as information continuously changes, nothing old remains the same, and what is currently being retrieved bears only a resemblance to what was previously sent. While this does not support the static view of memory, it does support the continuous integration of information, which is, after all, one of the functions in which memory plays a role.

4.1.2.2 Memory integration vs. segregation

Memory in motion is a balance of integration and separation tradeoffs (Dixon & Di Lollo, 1994). Firstly, to build up a memory leading to a unified percept of object continuity, one must integrate information across time and space. We shall examine such integration with the phenomena of apparent motion and visual persistencies. Of course, as objects move across space, the corresponding integration store must equally move or track.

Secondly, once a steady representation is built, the memory in motion can be used to detect and segregate rapid feature changes (Di Lollo, Hogben, & Dixon, 1994), including how to deal with onsets and offsets that may refer to the existing object or a new occluding one. Alternatively, a memory in motion can be used for smooth pursuit through an occlusion (Orban de Xivry, Missal, & Lefèvre, 2008).

Vision is not like a camera where all information is captured in one frame at one point in time. Instead, each portion of memory, a channel if you will (Aymoz & Viviani, 2004; Nemes, Whitaker, Heron, & McKeefry, 2011; Zeki, 2005), has its own integration, segregation and binding windows. Thus, when some external trigger requires a single

unified percept, these channels are not necessarily bound at the same point in time. The flash-lag effect will illustrate this phenomenon. Note that the channel concept does not work well for complex stimulus with motion (B. B. Lee, 2011).

4.1.2.3 Brief Stimuli for 'flash' memories

There are many forms of presenting brief stimuli, and these can exhibit flash memories of varying kinds. We will review the phenomenology here but note that our ecological case is when a part of a stimulus becomes

Illustration 85: Monkey swinging through trees
From (Crapse & Sommer, 2008).

occluded for a time, to reappear shortly as when a monkey swings through a tree. It is the occlusion that can cause the stimulus to flash as in the scenario depicted in Illustration 85 – a branch or leaf gets in the way of a scene; with lots of leaves and branches, scenes become mere glimpses or flashes of useful stimuli.

Our interest in these flash memories is in uncovering their minimum durations or time constants, as well as understanding their relation to motion and smear via compensatory/predictive mechanisms.

4.1.2.3.1 ISI and OSA

One purpose of this chapter is to investigate the performance of a flashed stimulus (under motion) after its offset, i.e., during the interstimulus interval (ISI) to the next flash or mask. These measures are shown in Illustration 86. However, as a first approximation, the
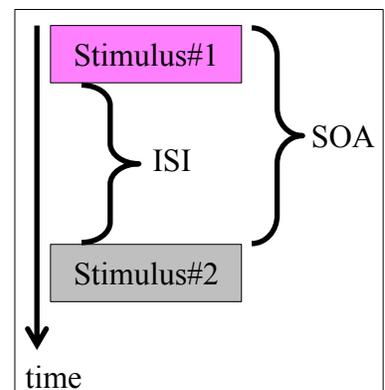
Illustration 86: ISI and OSA

time from the onset of one stimulus to another – the stimulus onset asynchrony (SOA) –

may be more important. For example, the identification of a stimulus lasting 240ms (with

ISI=0ms) is similar to a 13ms stimulus flash followed by 227ms blank ISI (Rucci &

Beck, 2005). Presumably, after 240ms, the stimulus has been transferred to VSTM and

any ISI afterwards does not improve performance. There is a performance departure at

40ms indicating some change in processing.

4.1.2.3.2   Flashed target

The brief flashing of a target (or fixation point) is a common cue to direct eye motion.

Often the target will disappear before the saccade completes. Nevertheless, the eye can be

positioned correctly despite loss of target and even with intervening saccades. This

requires a memory for the target *location*, but it need not be a *visual* memory, e.g.,

(Keith, Blohm, & Crawford, 2010).

In monkeys, such flashed cues are effective even when their duration is as short as 25ms

(Bruce & Goldberg, 1985) – flashed targets do not need to be recognized, merely acted

upon. Note that in this section the position information is maintained within the

sensorimotor *action* system. Position information can also be maintained in the

perceptual domain (Sapkota, Pardhan, & Van der Linde, 2011).

4.1.2.3.3   Flashed stimulus while dark adapted

If the lights within a dark room are briefly flashed, then a positive detailed afterimage

will appear that can last for as long as a minute (Pelz & Hayhoe, 1995). These

afterimages are stable in space despite eye movements. Indeed, if a small stationary light

remains present, it is the light that appears to move. With small afterimages, the opposite

occurs. These phenomena both suggest that eye movements can update an internal model

of the world. Note that the positive afterimage does not appear to conscious awareness

for one or two seconds after stimulus offset.

This paradigm requires a dark adapted eye and is not similar to our ecological case of

occlusion during photopic scene perception. The flash memory is extremely long

compared to the 250ms of iconic memories.

4.1.2.3.4   Gist

If a photograph is flashed for 20ms, then

ERP analysis indicates that it takes 150ms

to recognize whether the scene contains an

animal (Thorpe, Fize, & Marlot, 1996).

Manual reaction times to 'go' trials

averaged 445ms. Using a forced-choice

saccade task, however, reliable saccades to

the correct hemifield were initiated as

Illustration 87: Anatomical flow-chart for forced choice saccade task.
From (Kirchner & Thorpe, 2006), Fig. 9, showing anatomical connectivity with numbers indicating earliest response to flashed stimulus and average latencies.

quickly as 120ms implying an initial visual analysis of 95-100ms based on anatomical

latencies as depicted in Illustration 87 (Kirchner & Thorpe, 2006). Saccades to faces are

even faster and can start as early as 100ms (Crouzet, Kirchner, & Thorpe, 2010). While

these studies suggest that a *black box* memory trace of the stimulus need only last for

100ms, the anatomical connectivity suggests that every region operates on different

aspects of the stimulus and has access to their trace aspects for lesser amounts of time.

For example, in a partial-report paradigm that included *pseudo*-letters, participants reconstructed *actual* letters – the richness of perceptual experience is an illusion that "emerges from the interplay between partially accessible information and expectations" (De Gardelle, Sackur, & Kouider, 2009). Such expectations are likely distributed as tuning parameters throughout the visual system and do not represent a single visual memory store.

Animal recognition can even occur with a stimulus duration as short as 1ms, but this obviously requires a form of memory, possibly iconic and possibly due to recurrence (Thurgood, Whitfield, & Patterson, 2011). One way to reduce the effects of such memory is with a subsequent mask. Fei-Fei, Iyer, Koch, & Perona (2007) found that after a 27ms exposure followed by a mask, subjects could still identify low-level sensory-feature-related elements of a gist in their free-recall responses. There was an advantage to *animate objects*, vs. *animal* or *mammal* categories in terms of accuracy and frequency of report. Along with the view that gist is immune from inattentional blindness, this suggested consciousness awareness without the need of attention (Tsuchiya & Koch, 2009). More recent research indicates that even gist perception requires attention (Cohen, Alvarez, & Nakayama, 2011; Mack & Clarke, 2012). There is one problem with these gist studies. It is not clear that a mask eliminates all forms of memory-like processing, and it is possible that unconscious processing continues (Potter, 2012). This of course, highlights a problem with all perceptual studies, how much of the phenomenon to attribute to sensory processing, and how much to perceptual processing. One way to uncover the minimum sensory duration is via RSVP.

Using a rapid serial visual presentations (RSVP) protocol, a picture can be detected on the basis of meaning when presented with a duration of 13ms, i.e., at 80 pictures per second (Potter, 2012). In monkeys, 65% of the neurons that responded to faces retained their selectivity at presentation rates as fast as

Illustration 88: Rapid Serial Visual Presentation (RSVP) protocol
From (Potter, 2012), Fig. 2.

14ms/image (Keysers, Xiao, Földiák, & Perrett, 2001). The neural response did not indicate a fusion of images, although some integration in sequential images was likely. These are in line with the flicker fusion studies indicating a temporal period of 17ms Hecht & Shlaer (1936). It is of course, not clear what kind of memories persists. Potter suggested a conceptual short-term memory (CSTM).

Existing computational models may have trouble with such short presentations as they must actively clear their explicit or implicit memories. While ECM has a 10ms tick cycle for convenience, due to the priority of bottom-up information flows, it can currently handle stimulus changes every 10ms without interference. If a stimulus exists, ECM acts like it does not have "memory". If a stimulus is missing, ECM acts as if it does.

One should point out that the RSVP paradigm is not ecological. "The world has natural sequences and does not have a succession of random visual events" (Perrett, Xiao, Barraclough, Keysers, & Oram, 2009). A current stimulus sets up expectations for the next by temporal interactions that affect sensitivity and representational bias as shown in Illustration 89. This could allow for even faster recognition than suggested by standard RSVP. In line with this entire thesis, Perrett et al recommend that

Illustration 89: Tuning to isolated images and image sequences
From (Perrett et al., 2009), Fig. 9, showing diminished response to images in sequence (dotted line) compared to individual images (solid lines).

> *"notions of neural receptive fields and tuning for discrete stimuli may need reformulating to reflect a predictive function during the processing of more natural and continuously changing stimuli. We conclude that in a changing perceptual world, the values represented by cells' activities reflect a predicted future state rather than the present reality."*

RSVP is an example where there is no time for integrative processing and every frame represents a segregation activity. However, any system with persistence beyond 14ms would start to demonstrate *interference* between images – we can't use the terms *blur* or *smear* as these are defined with respect to motion, and *dissipation* seems inappropriate.

4.1.2.3.5 Iconic memory and visual persistence

The partial report paradigm was introduced by
Sperling (1960). He showed that in a whole
report, the *immediate memory span* for a list of
alphanumeric symbols was 4.5 items and that
this did not interact with an exposure duration
that varied from 15-500ms. If the stimuli were
arranged in rows containing no more than 4
characters each, e.g., (4/4/4) or (3/3/3), and if

**Illustration 90: Decay of available
information in (3/3/3) partial report.**
From Sperling (1960), Fig. 7.

the specific row could be cued by a low, medium or high auditory tone or light flash, then

this partial report paradigm indicated that all available letters could be identified.

However, this memory rapidly decays with cue delay as shown in Illustration 90.

Essentially the partial report advantage over the whole report is 250ms. Sperling

attributed this phenomenon to image persistence, but it was renamed *iconic* memory by

Neisser (1967) who wanted to avoid the many assumptions inherent in an *image* based

memory. However, even Neisser wondered whether the concept would be problematic

under motion as there would be overlap and interference among icons corresponding to

separates glances over the same scene (Haber, 1983).

> *"the conditions that produce iconic-like persistence do not exist under
> naturalistic perceiving, and if we had iconic storage of the previous instant of
> stimulation, perceptual chaos rather than continuity would prevail." – (Haber,
> 1985)*

After publishing more than 30 papers on (or involving) the phenomenology of iconic memory, Haber (1983) has called for the abandonment of the entire iconic memory concept.

Moreover, the duration of such memories seems to vary by psychophysical method and response criterion possibly measuring different entities (Breitmeyer, Kropfl, & Julesz, 1982; Coltheart, 1980; Loftus & Irwin, 1998; Long, 1980).

> *"Iconic memory in this neurophysiological perspective is thus a multilayered process, composed of a spectrum of levels of categorisation with different time courses."* – (Keysers, Xiao, Foldiak, & Perrett, 2005)

Caution needs to be exercised when translating the perceptual performance curves as used in iconic memory experiments to the degree of blur, smear and dissipation occurring within a specific neural locus, especially as memory seems to be distributed across loci.

4.1.2.3.6   Saccadic preview effect

If a pre-saccadic stimulus consistent with the target (a curve) is presented for a duration as short as 10ms, then it confers a preview benefit at post-saccadic processing (Khayat et al., 2004). This demonstrates that the entire curve is maintained in transsaccadic memory for the duration of the ~95ms saccade.

This paradigm is comparable to our ecological form.

4.1.2.4   Apparent motion

The percept of apparent motion can occur when an object is flashed sequentially at two different locations as shown in Illustration 91, and this may involve what we call 'flash' memories during the inter stimulus interval (ISI) between flashes.
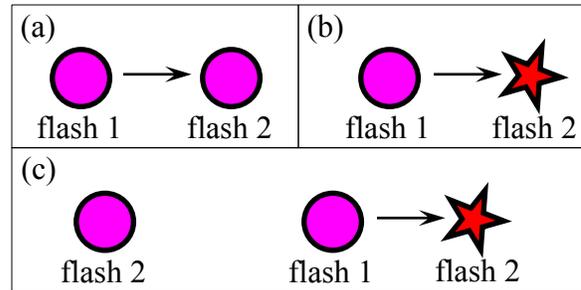
Illustration 91: Apparent motion.
After (Flombaum et al., 2009), Fig. 6.4: (a) arrow indicates apparent motion perceived between 1st and 2nd stationary flashes; (b) even if 2nd flash differs in some features; (c) and with distance trumping features.

Our ECM model does not currently support *direct* apparent motion, but we mention this paradigm for completeness, in part because it is within our ecological context – the ISI could represent an occlusion – and in part because it will inform the simpler paradigm modelled and tested within this chapter, i.e., relative real motion of stationary objects due to eye motion.

There is also a relation between apparent motion and the flash-lag effect. Namely, a single flash would have a long (~150ms) persistence, but two nearby flashes representing an object moving across frames would have a shorter persistence, ostensibly to reduce blur (Coltheart, 1980; Moore, Mordkoff, & Enns, 2007). But the lag still needs to be explained in relation to the apparent motion of object-like entities.

4.1.2.4.1   Object memory

Moore, Mordkoff, & Enns (2007) divide visual processing models into two classes, image based "painting" programs and object based "drawing" programs.

> *"Despite the appeal of image-based updating, **this alternative would carry considerable cost** [bolding ours] for later, higher-order processing. Because such*

updating would occur without regard to the organization of the scene in terms of surfaces and objects, many important distinctions would be lost with each resampling cycle. Image-based updating would fail, for example, to maintain region assignments to figure vs. ground, edge assignments to luminance change vs. surface orientation, as well as associations between discontinuous regions of a surface caused by occlusion."

While Moore et al. prematurely jump to an object based conclusion without any computational basis, e.g., the bolded text in the quote is a subjective intuition, we want to point out an alternative to their dichotomous thinking. ECM provides an intermediary conception. While ECM supports perception which could be object based, ECM is only a sensory image based system with generative capabilities. However, the larger (unimplemented) theory (p211) does segment the image into overlay (occlusion/transparency), fixation, and background layers so that all the given reasons to dismiss image-based models can themselves be dismissed. In particular, a visual scene becomes stable within the flow, so there are no issues with updating higher-level representations if based on a stable flow. Indeed, ECM suggests that motion can only be handled within such an underlying architecture.

Nevertheless, the phenomenology of apparent motion cannot be understood without the concept of an object-like representation – a higher level form of memory. When it comes to motion, data is highly contaminated by such theory.

4.1.2.4.2 Feature fusion

The apparent motion within sections (b) and (c) of Illustration 91 demonstrate *spatiotemporal priority* leading to an object persistence that may integrate or fuse

percepts across short glances (Flombaum, Scholl, & Santos, 2009). This may also explain the ubiquity of change blindness as our visual system is predisposed to detect objects and not feature changes. Whether visual information is merely sensed or consciously perceived, it is ultimately for adaptive *behaviour* and not for *seeing* (Goodale, 2008).

This smooth and continuous motion can also occur across saccades that include discrete objects transformations (Fracasso, Caramazza, & Melcher, 2010). In cases of apparent motion or transformational motion, there is no tracking of intermediates as shown in Illustration 92, so the smoothness of the percept must be a phenomenal construction after the second stimuli has been sensed. Indeed TMS pulses with SOA of 80-100ms seem to mask awareness and this is

Illustration 92: Transformational apparent motion across saccades
From (Fracasso et al., 2010), Fig. 1a) various transformational stimuli and the percepts evoked.

presumed to operate on feedback projections (De Graaf, Goebel, & Sack, 2011; Ro, Breitmeyer, Burton, Singhal, & Lane, 2003) and can also be demonstrated by object substitution masking (Di Lollo, Enns, & Rensink, 2000; Koivisto, 2012). However, neither TMS nor masking prevent unconscious location detection (Railo & Koivisto, 2012). Thus persistency effects are highly dependent on the specific neurobiological pathways used.

4.1.2.4.3 Tunnel effect

The ISI between the two flashes of apparent
motion are similar to the time period when an
object travels through a spatial gap such as a
tunnel with observers perceiving the hidden
movement (Burke, 1952). This apparent

**Illustration 93: Perceived trajectory of
apparent motion**
From (Burke, 1952), Fig. 3a. Configuration
of tunnel above, perceived motion trajectory
below.

movement will be perceived as 'S' shaped rather than a straight line when the trajectories

are not in line despite a total lack of visual evidence within the tunnel as shown in

Illustration 93. The fusion of features can also be found in the tunnel paradigm. If a

lemon rolls through a tunnel and a kiwi comes out at the appropriate time, a monkey will

take this as one object; if the timing is off, the monkey will take this as two (Flombaum,

Kundey, Santos, & Scholl, 2004).

4.1.2.5 Real motion

Visual processing requires a balance between integration and segregation (Dixon & Di

Lollo, 1994). With flashes, discrete stimuli are integrated over time and can provide

percepts such as apparent motion. Once integrated into a stable and continuous percept,

changes can be detected and segregated. In this section we start off with real integrated

motion and review phenomenology associated with sudden changes such as nearby

flashes, or occlusions.

Originally, eye fixations were considered as equivalent to a discrete stimulus or flash, and

saccades corresponded to an ISI where information was suppressed. However, now we

know that vision undergoes continuous motion where drift, jitter and microsaccades are some of the continuously occurring fixational eye movements (Martinez-Conde, Macknik, & Hubel, 2004; Murakami, 2006; Rucci, 2008). While experimental paradigms are rife with discrete stimuli where smear is not a problem, there are few ecological examples of discrete stimuli. Instead, stimuli should always be considered as under continuous motion (Haber, 1985; Turvey, 1977), and smear is a perennial problem to be managed.

During real motion, either due to object motion, or eye motion, we are interested in the maintenance of stimulus during the ISI, and the prevention of smear.

4.1.2.5.1   Occlusion processing

The problem with occlusions is that they can disrupt integration, and they are ambiguous. If the stimulus offset event represent the temporary onset of an occlusion (to be followed by the occluder offset), then the object could be maintained until the occlusion disappears. If the stimulus offset event represents the permanent disappearance of an object, then in this case there is no occlusion onset but an onset of the background. Also, object memory could be erased, however, how can one consciously determine which object permanently disappeared unless there is some object persistence? If the stimulus offset event represents the permanent appearance of an occluder, then old object information could be discarded and a new object created for the occluder. Again, how to account for knowing which object disappeared? As onset, offset, foreground, background, occluder and object processing all seem to take place, it is no wonder that

numerous lower and higher cortical areas are involved in the active spatiotemporal

maintenance of dynamic objects through occlusions (Shuwairi, Curtis, & Johnson, 2007).

4.1.2.5.2 Smooth pursuit through occlusion

The eyes cannot move smoothly without a

smoothly moving stimulus, but when the stimulus

becomes occluded, smooth pursuit continues by

maintaining a predictive dynamic representation.

The relation between smooth pursuit, memory in

motion and blur will be reviewed in this section.

Smooth pursuit is one mechanism to prevent

target blur due to relative retinal motion at the

expense of background blur. Indeed visual

sensitivity is enhanced during smooth pursuit, and

while this might facilitate tracking, it may also

help to reduce blur of the background (Schütz,

Braun, & Gegenfurtner, 2009), e.g., enabling a

faster "shutter speed". Nevertheless, the

background is much less smeared than expected,

and this may be due to compensatory (or

predictive) mechanisms based on extraretinal eye-

movement information (Bedell et al., 2004).
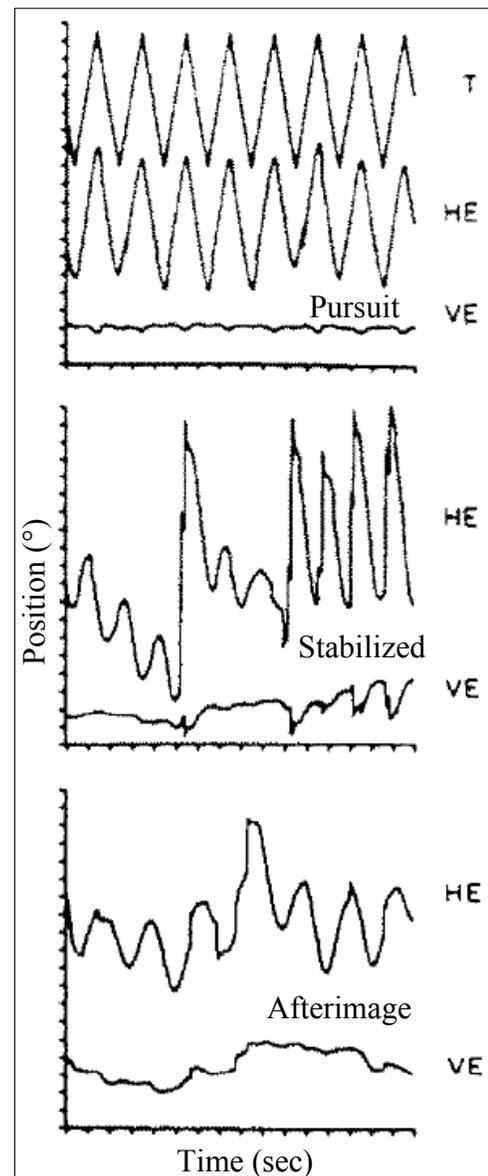
Illustration 94: Smooth pursuit requires stimulus
After (Cushman, Tangney, Steinman, & Ferguson, 1984), Fig. 1: (Top) horizontally moving target (T) being smoothly pursued by horizontal eye movement (HE) and non-vertical eye movement (VE); (Middle) showing irregular eye movement when target electronically stabilized; (Bottom) showing pursuit of afterimage.

Smooth eye motion cannot be made in the absence of a smoothly moving stimulus (Lisberger & Westbrook, 1985; Yarbus, 1967) as shown in both the horizontal and especially the vertical eye movement traces in Illustration 94, Illustration 95, and Illustration 78 on page 98. However, tracking a generated stimuli, such as an afterimage can be smooth (Heywood & Churcher, 1971), even though it will not track accurately for lack of correct feedback. An afterimage is just one of
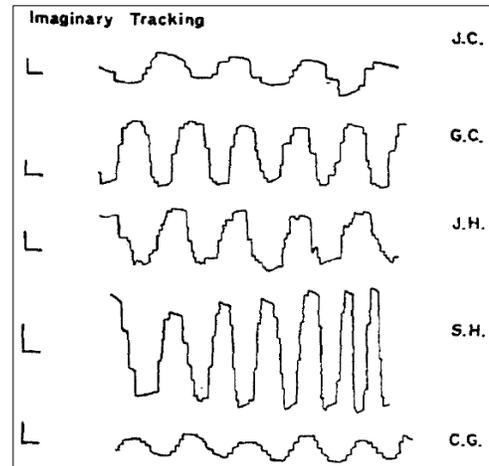
Illustration 95: Non-smooth tracking of imaginary pendulum in the dark
From (Heywood & Churcher, 1971), Fig.1: Representative EOG tracings of 5 subjects; Calibration mark is 20° vertical by 1.5s horizontal.

several perceptual, but not sensory stimuli that are involved with smooth pursuit. We will review many others.

Smooth pursuit does not track salient sensory stimuli such as points or edges, but the motion of perceptual objects which may take 90ms to determine (Masson & Stone, 2002). Indeed, the overtly tracked point could be invisible to the sensory system, and instead, could be a perceptual construction. For example, it ordinarily tracks perceptual constructions such as the invisible centre of an object. It can track the completed contours extrapolated from an object in the periphery (Steinbach, 1976) or the midway between two points (Wyatt, Pola, Fortune, & Posner, 1994). In anorthoscopic viewing (as in Leibovitz, p201), pursuit increased after a figure percept was reported indicating that the percept came first, enabling the pursuit, and it was not the pursuit that enabled the image to be painted onto the retina and thereafter forming a percept (Fendrich, Rieger, &

Heinze, 2005). Smooth pursuit is 4% faster when coordinated by arm movements (Snyder, Calton, Dickinson, & Lawrence, 2002). Such coordination has a neural locus in area LIP of the posterior parietal cortex (PPC) (Hagan, Dean, & Pesaran, 2012).

Smooth pursuit occurs in apparent motion when image makes discrete jumps and ISI < 150ms, but becomes rough pursuit if ISI>150ms (Morgan & Turnbull, 1978) and contains interrupting saccades. However, the perception of smoothness has a different cutoff than the action of smooth pursuit.

Conversely, there are cases where an unconscious/unperceived visual signal is tracked (Spering & Carrasco, 2012; Spering, Pomplun, & Carrasco, 2011). This may reflect the vision-for-action pathways rather than the vision-for-perception-then-action pathway.

When the target of a smooth pursuit becomes occluded or blanked, then the eyes will reduce speed to zero within 100-200ms (Barnes & Collins, 2008a; Mitrani & Dimitrov, 1978; Pola & Wyatt, 1997). However, if the target is expected to reappear, then eye velocity will recover to a scaled prediction of target velocity lasting as long as 4s (Becker & Fuchs, 1985; S. J. Bennett & Barnes, 2004). This predictive smooth pursuit will be interspersed with several "catch-up" saccades to correct for predicted target position (Orban de
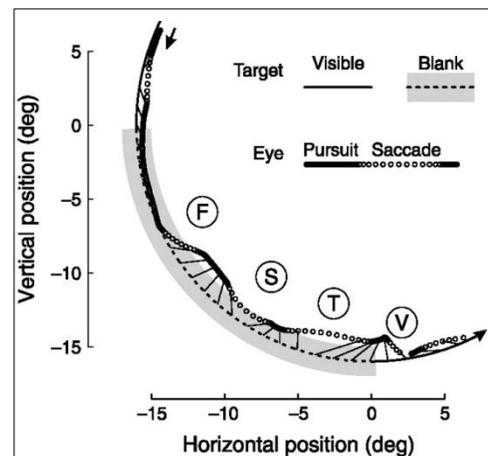
Illustration 96: Predictive smooth pursuit and corrective saccades during occlusion From (Orban de Xivry et al., 2009), Fig. 2: Pursuit interspersed with (F)irst (S)econd and (T)hird predictive saccades followed by a (V)isually guided saccade when target reappears.

Xivry, Bennett, Lefèvre, & Barnes, 2006) as shown in Illustration 96. Unlike corrective saccades during *visual* smooth pursuit, the goal of corrective saccades during *predictive* smooth pursuit seems to be to jump ahead of the target so as to minimize position error upon the object's eventual reappearance (Orban de Xivry, Missal, & Lefèvre, 2009). Predictive information can be accumulated within 150ms prior to occlusion, and must thereafter be held in short term memory to drive predictive behaviour throughout the blank period (Barnes & Collins, 2008b).

Anticipatory smooth pursuit eye movements can be initiated ~200ms before onset of stimulus motion and will be unaffected by visual stimulus feedback for ~80ms. The neural locus for these expectations can be found in a population of neurons in the supplementary eye fields (SEF) (De Hemptinne, Lefèvre, & Missal, 2008).

The expectations could involve *imaginary* objects. When monkeys were trained to saccade to imaginary objects cued peripherally (Wyatt et al., 1994), some neurons in MST were activated for both actual dot targets and imaginary ones, while MT neurons were only activated by a real target (Ilg & Thier, 2003). In any case, during an occlusion, smooth pursuit is driven via a dynamic representation of target motion in short-term memory (Orban de Xivry et al., 2008) – what we refer to as memory in motion.

For optimal performance under noisy environment, spatiotemporal integration must still occur for smooth pursuit. A temporal filter with a 25ms time constant accounts for 50% of the variance in eye motion (Osborne & Lisberger, 2009).

4.1.2.5.3   Object creation

If a flash occurs at the same position as a continuously moving object, then the flash location will be perceived correctly, but the object will appear as advanced. This is shown in the top row of Illustration 97. For some reason, this phenomenon was

Illustration 97: Object creation
From (Nishida, 2011), Fig. 4: (Top) continuous object perceived ahead of flash; (Bottom) abruptly changed object perceived as new and aligned with flash while old object still advanced (Moore & Enns, 2004).

named in reverse, the *flash-lag effect*, and will be detailed shortly. However, if the object changes size during the flash (bottom row), it will be perceived as a new object correctly positioned, while the old object still appears in an advanced position (Moore & Enns, 2004). In the case of apparent motion, significant changes will be fused into one object, but during continuous motion, feature changes that are substantial, abrupt and transient will disrupt object continuity (Moore & Enns, 2004) – they will be segregated into a new object. Such new objects will capture attention (Rauschenberger, 2003), but take longer to process (Kanai, Carlson, Verstraten, & Walsh, 2009) and hence appear to lag. Thus the final percept in Illustration 97 shows both objects when there was never more than one. Apparently, the binding of all objects into a "unified percept" does not occur at the same time, but this concept of unified perception invokes the snapshot metaphor of perception which is problematic. Binding is a slow (Holcombe, 2009) attribute dependent process (Bartels & Zeki, 2006).

It should be mentioned that in general new objects capture attention only if there are attentional resources available and if the old objects are actually being attended. If the old

objects are static they may be ignored, lose their VSTM representation or "oldness" and new objects may not be distinguishable from all the old ones, as they too are considered as equally new. However, if the old objects are moving, they will be maintained in VSTM during a momentary (~10ms) occlusion enough to retain attention, so that a new object inserted during that occlusion will be noticed as the only new entity (Chua, 2011). The occlusion, in this case, is to mask the sensory transient that could signal a new object (Hollingworth, Simons, & Franconeri, 2010). Note that in an ecological environment, there are innumerable new objects all the time, and VSTM cannot maintain all the old objects, hence change blindness.

4.1.2.5.4   Multiple object tracking (MOT)

The multiple object tracking (MOT) paradigm (Pylyshyn & Storm, 1988) is used to study the interaction between sensory and perceptual systems under the influence of attention as shown in
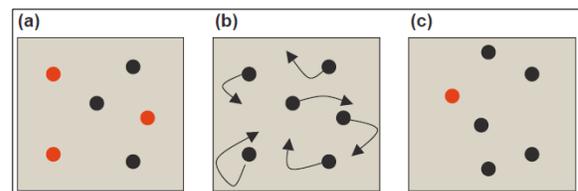
Illustration 98: Multiple object tracking (MOT)
From (Cavanagh & Alvarez, 2005), Fig. 1: (a) targets to be tracked briefly identified; (b) targets tracked among distractors; (c) partial report test on random object to determine if it was a target

Illustration 98. Spatial attention can be divided so as to track 4-5 objects independently. However, this may be a hemifield limited resource so the actual limits may be more like two objects per hemifield (Alvarez & Cavanagh, 2005). Also see (Delvenne & Holt, 2012). At slow speeds, 8 objects can be tracked and at fast speeds only 1 (Alvarez & Franconeri, 2007). However, spatial attention is affected by crowding, surround suppression (Petrov, Popple, & McKee, 2007) and ensemble coding (Bulakowski, Post, & Whitney, 2011). By keeping objects small and gaps large (Herrmann, Montaser-Kouhsari, Carrasco, & Heeger, 2010), it was found that a practically unlimited number of

objects can be tracked suggesting a low-level parallel processing system with local

systems tracking individual objects and inhibiting nearby ones (Franconeri, Jonathan, &

Scimeca, 2010). This is also consistent with a top-down and bottom-up competitive

interaction for segregating a scene into candidate objects (Beck & Kastner, 2009; Keitel,

Andersen, Quigley, & Müller, 2012; McMains & Kastner, 2011). In many cases, the

number of items is traded off against precision. For example, it was found that the

direction information in MOT is a continuously divisible resource and there is no fixed

bound on number of objects (Horowitz & Cohen, 2010). Similar findings of unlimited

object capacity exist in other paradigms (Bays, Catalao, & Husain, 2009; Wilken & Ma,

2004) suggesting an analytical move away from items and towards structured

representations (Brady et al., 2011). Indeed, there may be a fractionation of memories

corresponding to the perceptual/sensory divide such that

> *"observers could not increase the number of representations in VWM by reducing*
> *the precision of the representations… In contrast, we found that observers could*
> *trade the precision of representations for storage of a greater number of items at*
> *an **earlier stage of representation**" – (Zhang & Luck, 2011)*

A FINST is a proposed "mechanism" that would track an identified object thereby

performing a "preconceptual reference function" (Pylyshyn, 2001). However, no

computational account has ever been provided, so it becomes impossible to ascertain

when merely preconceptual/sensory processing occurs, and when tracking occurs. In one

sense, the hard problem of binding has been hidden within the FINST. Similarly, the

integration of sensory information that would facilitate the construction of the next

FINST is undetermined. A FINST is therefore nothing more than a surrogate for a theory (Gigerenzer, 1998) – it is certainly not a mechanism.

One of the functional requirements for a FINST is for tracking identified objects in a preattentive, automatic and data driven manner. It is not clear what the data (features or properties) are that enable FINST tracking as observers fail to detect changes in colour or shape when an object passes behind an occluder. Therefore a FINST tracks and *maintaining identity* "operationalizes the notion of `primitive visual object' as whatever allows preconceptual selection and MOT" (Pylyshyn, 2001). When parts of an object are connected, studies indicate that attention is allocated to higher level objects and that features must be merged (Scholl, Pylyshyn, & Feldman, 2001), including the suppressive effects of features associated with distractors. But this data fusion cannot occur within an automatic data driven system, and Pylyshyn is especially silent on how this could possibly be handled in the conception of FINST. One puzzling result was that while objects could be successfully tracked, their identities were poorly recalled (Pylyshyn, 2004) and this negates a primary purpose for having FINSTs. It has therefore been suggested that MOT is nothing more than attention in action (Scholl, 2008).

Regardless of the status of MOT and FINST within cognitive theories, the MOT phenomenon and experimental paradigms have provided a wealth of information on tracking, and with this comes with some paradoxes. For example, while purpose of a FINST is to track objects to support MOT, one would think that this would include the ability to do so behind momentary occluders by extrapolating object trajectories. Instead,

objects are better tracked when they reappear at the old location rather than at their

expected location (Franconeri, Pylyshyn, & Scholl, 2012; Keane & Pylyshyn, 2006).

Note that MOT is a task level phenomenon and has not been related to sensory

phenomena of real or apparent motion percepts. Moreover, the objects had disappeared

behind their occluders for periods longer than sensory or iconic memories.

4.1.2.6 Flash-lag effect

In the flash-lag effect (FLE) (Hazelhoff &

Wiersma, 1924; MacKay, 1958; Nijhawan, 1994),

the position of a flashed item will appear to lag

behind a moving yet unattended object as shown

in a popular configuration within Illustration 99.

Attending to the moving object will eliminate the

Illustration 99: Flash-lag effect
From (Nijhawan, 2002), Fig. 1. (a) When fixating on the moving black ring, an internal flash is perceived properly; (b) when fixating elsewhere, the internal flash appears to lag behind the ring.

lag (Namba & Baldo, 2004). This effect is paradoxical as the cognitive processing of a

flashed stimulus is considered to be much faster than that of a moving stimulus

(Nijhawan, 2008; Walker & Irion, 1982), although some presume opposite delays

(Arnold, Ong, & Roseboom, 2009). This differential processing effect perhaps applies to

the cases of stimulus onset and offset as well. In a related phenomenon, the duration of a

brief flash is perceived as longer when attended (Enns, Brehaut, & Shore, 1999;

Yeshurun & Marom, 2008). Both of these effects appear to be modulated by top-down

attentional influences.

Such differential processing, both in using attentional mechanisms, and in using multiple

speeds in processing pathways is not considered by most visual models nor the specific

models in this chapter as we are primarily interested in the dissipation dynamics of memories after stimulus offset under continuous eye motion as they relate to smear. These effects illustrate the woeful inadequacy of most visual models that are primarily a single bottom-up feedforward hierarchy. However, even explanations for the flash-lag effect are woefully inadequate as they rest on non-empirical intuitions, assumptions, speculations, and mechanistic conjectures. Precise theoretical (or computational) models of mechanism are essential (Trautteur, Datteri, & Santoro, 2008). Such overgeneralized yet underpowered hypotheses are examined next.

Whether pathways have differential internal delays or not, there are always delays of some kind such that by the time we perceive the location of a moving object, it has already moved on. The sensorimotor system can compensate for this via a predictive control mechanism (Davidson & Wolpert, 2005; Nijhawan & Wu, 2009), and it is hypothesized that visual prediction also occurs for moving stimuli (Nijhawan, 2008). The flash is a discrete stimulus, apparently stationary, therefore visual motion prediction is not applied, and hence it appears to lag. During smooth pursuit of the object, both the tracked object and the flash are effectively stationary, and so no lag is perceived. Predictive mechanisms can be found throughout the brain (Bubic, Von Cramon, & Schubotz, 2010) and are also found throughout the ECM architecture. In particular, the persistence of the stimulus after its offset is maintained by a mechanism that predicts the visual information flow due to eye motion.

An alternative explanation for the flash-lag effect involves motion blur with direct relevance to this chapter. It is hypothesized that blur processing reduces persistence of moving objects in relation to flashed objects (Nishida, 2011) so that binding occurs after the flash is over with a persistent trace of the flashed stimuli – it is this very trace that we model in this chapter. Indeed motion has been found to induce mislocalization of blurred images (Y. X. Fu, Shen, & Dan, 2001). Alternatively, the longer persistence of the flash may be due to the inverse relation to the stimulus duration (Walker & Irion, 1982). When subjects were induced to notice the flash onset, the FLE disappeared, implying that it is indeed the persistent trace of the flash that is ordinarily bound (Gauch & Kerzel, 2009).

Possibly, the object creation explanation mentioned earlier can suffice without the need for prediction – the flash is a new object and is simply bound later (Nishida, 2011).

4.1.2.7   Maintained/generative memory

We liken the removal of a stimulus as due to an occlusion, with the visual system maintaining and tracking the objects until they reappear. Thus, the length of the occlusion equals an interstimulus interval (ISI) and the duration of the memory we wish to analyze. The fact that the stimulus never reappears in the tests is irrelevant. In this section, we review the phenomenology of such memory maintenance, generation, and accounting for motion.

4.1.2.7.1 Multiple object tracking through occlusions

Multiple object tracking (MOT) through occlusions and the perception of objects as persisting through such momentary occlusions may feel effortless, but Flombaum, Scholl, & Pylyshyn (2008)
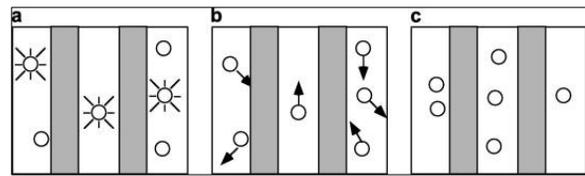
Illustration 100: Multiple object tracking through occlusions
From (Flombaum et al., 2008) Fig. 1: a) targets distinguished from distractors; b) targets tracked while the detection of random probes indexed attention; c) targets identified.

found that attention was directed at both targets and distractors when occluded and that probe detection was worse when these objects where visible. In these tests, attention was studied rather than the duration of object memories. While object based attention may be directed to targets and inhibited to distractors, this did not happen during occlusion. This suggests that occlusion processing may be at a lower, *sensory* level despite the need for an "attentional high-beam" to *perceive* the results.

Possible neural correlates to the inference of motion for moving targets in occluding trials, as compared to stationary targets in 'blink' trials, was found in the primate posterior parietal cortex (Assad & Maunsell, 1995). In these tests there was no extra retinal information as the eyes did not move, and indeed no sensory information during the blink or occlusion. The paradigm will be slightly different than in our tests as instead of inferring motion for moving objects, we use the extra retinal knowledge of eye motion over stationary stimuli to infer relative motion. While the mechanism to maintain object position may differ, the point is that such positions are indeed maintained across occlusions.

4.1.3   Stabilized images

The problem of image smear is related to the problem of stabilized images. Obviously any system that managed to keep images stable despite eye motion would not have image smear in the first place, but then the problem of image smear is passed onto the problem of maintaining image stability. While an endpoint in visual processing might be a stabilized image in spatiotopic coordinates, this does not negate that fact that prior neural systems do have memories and appear to be retinotopic. How can one avoid smear there?

4.1.4   Stephen Grossberg

In Test 4 we will be comparing ECM to the set of computational models developed by Stephen Grossberg and his many associates for several reasons. Firstly, the Grossberg set of models are the most complex and have attempted to deal with motion. Secondly, Grossberg has a similar desire to unify multiple phenomena within one model. Out of all the existing computational modeling systems, the Grossberg set of models represents the closest system to ours in spirit. Of course, they have thirty years of development behind them. This section serves as background information for the discussion in Test 4.

4.1.4.1   Unity and coherence, but without motion

In 1983, Grossberg argued for a coherent and unified computational theory of visual processing that would take local and global interactions between depth, length, lightness and form perception into account. It is important to note that motion perception was not considered a foundational aspect, but perhaps this is due to the fact that there was little knowledge of motion at the time. This specific oversight would ensure that Grossberg's models, like all before and after, would be *neurocentric*. However, like ECM, biological plausibility, recurrence and dynamics were recognized as important aspects of

architecture. Unlike ECM, Grossberg wanted to account for perceptual phenomena, while ECM uses them to induce a sensory system.

4.1.4.2 BCS, FCS & consciousness

In 1984, Cohen & Grossberg published a theory of dynamic brightness perception that included a boundary-contour system (BCS) and a feature-contour system (FCS) as shown in Illustration 101. The BCS and FCS processes would remain an integral aspect of all of Grossberg's subsequent models and can be thought of as
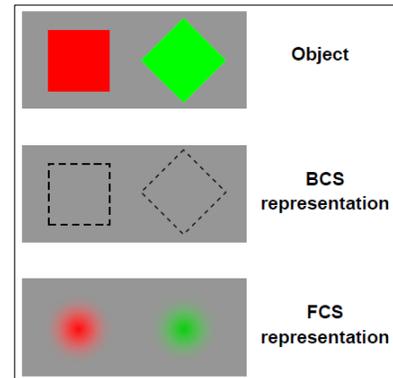
Illustration 101: BCS and FCS pre-interaction representation. From (Breitmeyer & Tapia, 2011),

the basic units of his architecture. BCS would be responsible for detecting and completing borders, while FCS would fill-in the surface features defined by the BCS via a diffusive mechanism. It is only subsequently that the interaction of these *invisible* activity patterns would enter perceptual conscious awareness. That is why the border representation within BCS is shown via a dotted line in – it is currently pre-attentive.

In "2   Chapter:" (p9) and "Appendix B:" (p231) we show that ECM has emergent behaviour that acts as if it has both a BCS and FCS, but without such complexity. Moreover, ECM has no explicit filling-in process and no diffusion. Finally, ECM explicitly avoids explaining perception or consciousness, but surprisingly, it does imply that the locus of consciousness is "in the flow".

138

4.1.4.3 Colour

Grossberg's basic ideas were further refined in (Grossberg & Mingolla, 1985a) where colour was first mentioned as one of the features that FCS could fill-in, but all examples remained in black & white, and there was no explicit neural support for colour. The complementary colour of some illusions was first discussed, but not demonstrated, in relation to inhibitory processing within some forms of BCS stimuli arrangements (Grossberg & Mingolla, 1985b). There was no colour-opponent processing, and in this, there are similarities with ECM where colour is processed as separate LMS channels without opponency. ECM does however demonstrate all results in colour but these are mere indexes to sensory processing that can support perception. ECM does not yet have an account of the negativity of colour since that is an aspect of perception.

4.1.4.4 FACADE

In (Grossberg, 1991, 1994), form-and-colour-and-depth (FACADE) are united, where form is supplied by BCS, colour is a feature supplied by FCS and depth is new. Depth provides for functional similarities to the segmentation layers of Leibovitz (p201). However in Grossberg's approach, perception is much more direct with less concern for visual dynamics. If an object is completely occluded it fades away immediately. Within ECM, the occlusion would be automatically routed to an overlay/occlusion layer, while the hidden object remains in the attentional layer where it is maintained and amenable to spatiotemporal integration.

In all the tests within this chapter, the entire stimulus is removed after a brief presentation. In a non-layered system such as FACADE, this would be considered as a

stimulus offset. In a layered system such as ECM, this is considered as the onset of an occlusion to be ignored for the short term.

4.1.4.5   Motion detection

In (Grossberg, 1987) the possibility of using motion sensitive cells was first mentioned, but the *motion BCS*, as distinct from the *static BCS*, was not described until (Grossberg & Rudd, 1989, 1992). It is important to note that motion information is not used to spatially maintain visual objects through occlusions. Instead, the motion BCS improves segmentation, can resolve motion ambiguity (the aperture problem), capture attention, and provide *formotion* percepts to conscious awareness (Baloch & Grossberg, 1997; Grossberg, Léveillé, & Versace, 2011).

Most simulation results from the Grossberg models do not take real object motion into account. We suspect that this is because the problem of smear (described next) would rear its ugly head. Most phenomena are related to apparent

Illustration 102: Inter-frame apparent motion. From (Baloch & Grossberg, 1997), Fig. 10c) arrow showing induced/apparent motion.

motion where smear is not a problem when an object is shown greatly displaced on two image frames, e.g., Illustration 102.

4.1.4.6   Smear

> *"The positive feedback can, however, maintain itself for a long time after visual inputs terminate. Thus, the very existence of cooperative linking could seriously*

degrade perception by maintaining long-lasting positive afterimages, or

smearing, of every percept.” – (Grossberg, 1991)

FACADE maintains information via resonance among neurons of retinotopic reference frame. Upon motion, this invariable leads to image smear as shown in Illustration 117 (p149). Grossberg recognized this problem in 1991, and suggested an inhibitory process to reset such resonance. Further discussion occurs in (Grossberg & Rudd, 1992).

4.1.4.7   Smear reduction via reset

One way to reduce smear is to detect stimulus offset and use that as a *reset* signal to inhibit the maintenance of old borders (Francis, Grossberg, & Mingolla, 1994; Francis & Grossberg, 1996; Grossberg & Grunewald, 2002). This mechanism will be further analyzed in Test 4.

4.2   Emergic Model

In order to explain the phenomena of this chapter, we have taken the ECM architecture as summarized on page 5, and made no changes. For the flash memories under motion we have two intuitions as to how ECM might explain the results.

Firstly, upon stimulus offset, bottom-up information, such as the blackness of a blink, or the whiteness of a blank screen would be routed to a future implemented overlay layer where it can be ignored (see p211). That is why this simulation can get away without such an implementation. Effectively, this suppresses information from going upwards into the first level RFs of the fixational layer.

Secondly, when bottom-up information is not available, the *handle missing data* function will simply make use of lateral information, and these form a stimulus memory of sorts. If there was no motion, the lateral information flows are equivalent to perfect recurrence – each RF would send its output back to itself – and such memory could be engineered to last forever without decay. However, there is continuous jitter, and in these tests, the eye does undergo a smooth change in fixation.

Thirdly, therefore, the *maintain information coherence* function shifts information to account for planned eye movement (without recurrence). For example, if the eye is currently planned to move exactly 2.5′ left that is equivalent to having the visual scene shift exactly 2.5′ to the right, so all the current visual information is shifted laterally to the right by 2.5′ to predict the actual visual input. The purpose of shifting is not prediction per se, but to align information, whether new stimulus or old memories so that they cohere within an RF.

Shift compensation can never be in integral RF units throughout, as RF sizes change by eccentricity, and their positions are heterogeneous. Thus an RF can be thought of as sending its content laterally, and each destination RF will pick up a portion of the content based on the amount of overlap between its spatial extent and that of the visual content. This leads to a form of memory dissipation that depends on the stimulus characteristics, but without the addition of blur or smear. As such, this is an example of a continuously changing memory in motion – a *flowcentric* approach to cognition.

4.3 Test 1-4: Robust stimuli

In these tests we characterize the memory behaviour of the ECM architecture after a briefly presented stimulus has been removed and while the eye undergoes moderate motion. To demonstrate the robustness of this spatiotemporal filling-in phenomenon, the model was examined across a wide range of stimuli.

4.3.1 Methods

The methods used are identical to those of page 80 except as indicated herein.

4.3.1.1 Subject analysis summary

This virtual agent was stochastically generated based on the development parameters detailed in the supplementary material (s115) while its photoreceptor mosaic is shown in Illustration 103. The visual angle for the entire photoreceptor mosaic was 88.51′ while the visual angle for the 1st, 2nd and 3rd RF levels was 98.86′, 104.18′ and 110.59′ respectively.
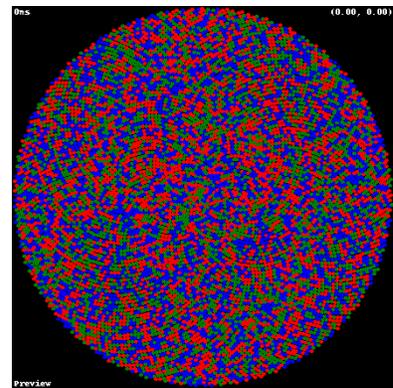
Illustration 103: Photoreceptor mosaic of agent used in "flash memory" experiments 1-4

Originally, stimulus offset was simulated within ESS by manually detaching the *person* from the *world view*. However, this was simplified by a set of blink related behavioural parameters and thus becomes a detail of *person* rather than procedure.

4.3.1.2   Stimuli

Similar stimuli as in tests 1-3 of Leibovitz (p60) with an additional square were used and are shown as viewed by *person* in Illustration 104, Illustration 105, Illustration 106 and Illustration 107. They are fully described in the supplementary material (s116).
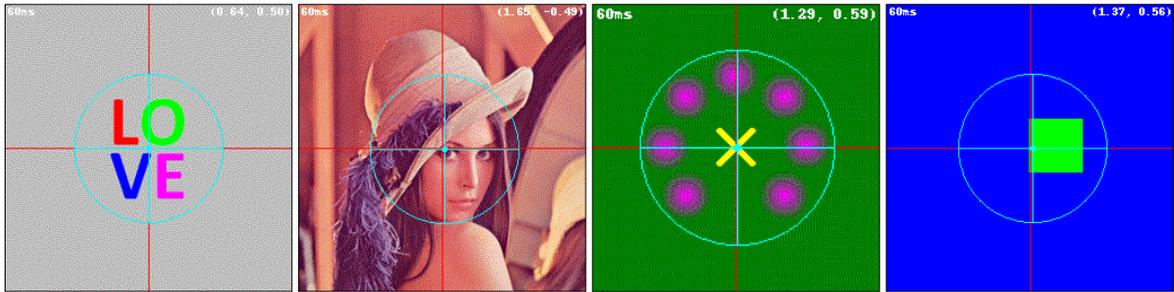
Illustration 104: Test 1 stimulus (Love)   **Illustration 105: Test 2 stimulus (Lena)**   **Illustration 106: Test 3 stimulus (Chaser)**   **Illustration 107: Test 4 stimulus (Square)**

4.3.1.3   Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to move its eye in the following pattern.

- Move right for 37 ticks at a speed of 1′/tick (± jitter)

The artificial brain was monitored from tick 6 thru 37 as it takes 6 ticks before motor plans take effect.

The stimulus remained on for only 6 ticks. It can be considered as a flashed stimulus or a brief presentation. The only reason that the stimulus duration was 6 ticks rather than 1 is so that stimulus offset coincided with the start of eye-movement.

Originally, stimulus offset was simulated within ESS by manually detaching the *person* from the *world view*. However, this was simplified by a set of blink related behavioural parameters and thus becomes a detail of *person* rather than procedure.

4.3.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s119, s130, s141, s152). When presented as an animation (Leibovitz, 2012a), all the results for this chapter can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=210
```

In this section, we summarize and discuss various highlights. We also extract memory cohesion times.

4.3.2.1   Behaviour

The dynamic behaviour of spatiotemporal filling-in for memory under moderate motion is shown for the 1st level RFs in Illustration 108, Illustration 109, Illustration 110 and Illustration 111. Similar behaviour occurs at the 2nd and 3rd level RFs. Labels over a grey background represent occurrences of eye movement. Labels in red represent "blink" periods without stimulus to the photoreceptors. Therefore the level 1 receptive fields (RFs) are impacted one tick afterwards. The stimulus is flashed on for the first 6 ticks so that eye motion start can coincide with stimulus offset. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
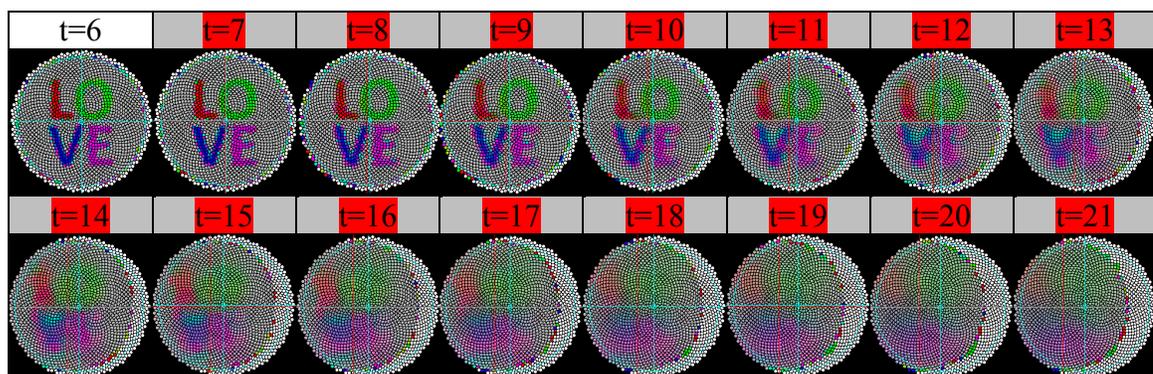
Illustration 108: Memory-in-motion filling-in behaviour for Love
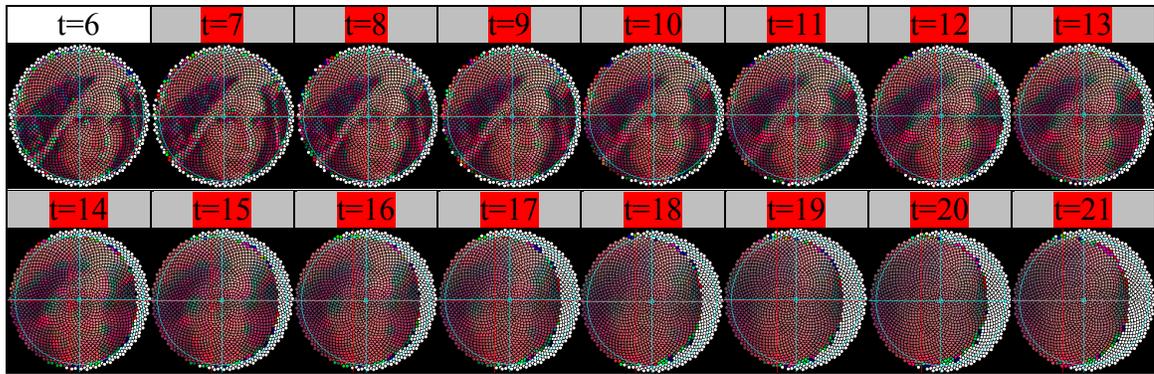
Illustration 109: Memory-in-motion filling-in behaviour for Lena
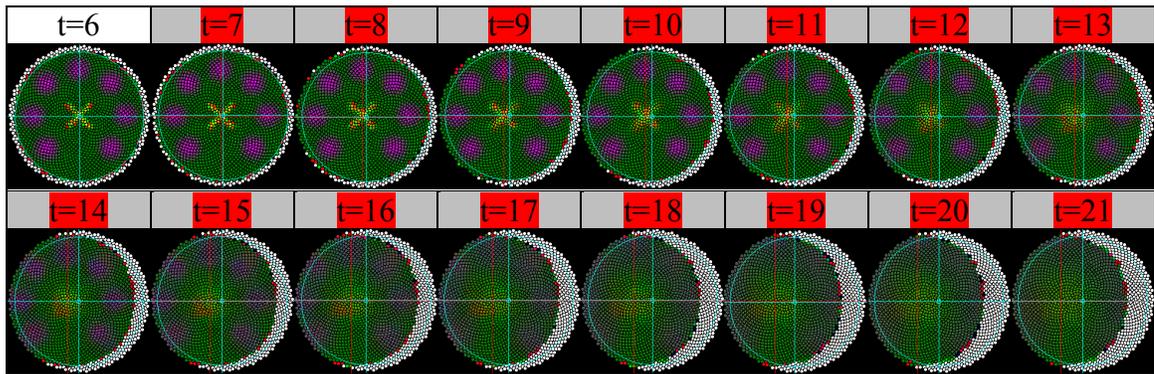
Illustration 110: Memory-in-motion filling-in behaviour for Chaser
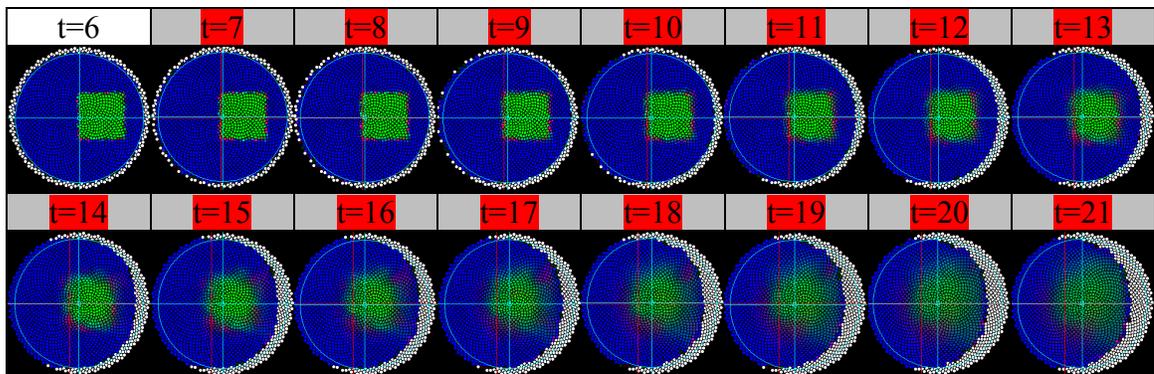
Illustration 111: Memory-in-motion filling-in behaviour for Square

4.3.2.2   Performance

For test 1, the gist of the stimulus is subjectively considered to have lasted till tick 14. The letters 'O' and 'V' are still
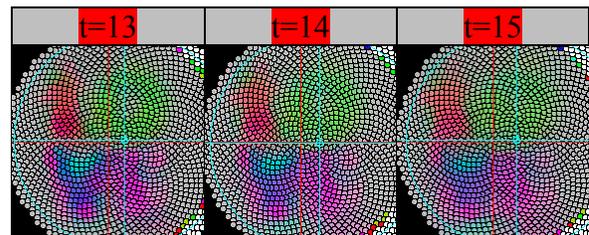
Illustration 112: Performance with Love

discernible; however the blue letter 'V' appears to have been separated into lilac, blue

and cyan components. The letter 'L' is still vertical but has lost its footing. The letter 'E' is ambiguous but retains colour, placement and aspects of shape. After tick 14, the gist is considered to have dissipated into the local average scene statistics as shown for the 1st level RFs in Illustration 112. This provides a memory in motion lasting 7 ticks (from 8 thru 14). Similar examples can be found for the 2nd and 3rd level RFs.

For test 2, the gist of the scene is subjectively considered to have dissipated into the local average scene statistics after tick 14 as shown for the 1st level RFs in
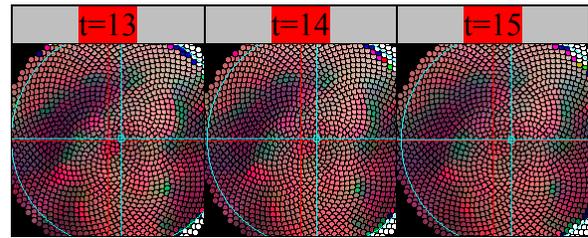
Illustration 113: Performance with Lena

Illustration 113. This provides a memory lasting 7 ticks (from 8 thru 14). Similar examples can be found for the 2nd and 3rd level RFs.

For test 3, the gist of the scene is subjectively considered to have dissipated into the local average scene statistics after tick 19 as shown for the 1st level RFs in
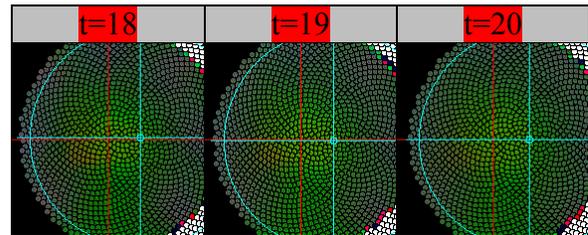
Illustration 114: Performance with Chaser

Illustration 114. This provides a memory lasting 12 ticks (from 8 thru 19). Similar examples can be found for the 2nd and 3rd level RFs.

For test 4, the gist of the scene is subjectively considered to have dissipated into the local average scene statistics after tick 18 as shown in Illustration 115. This
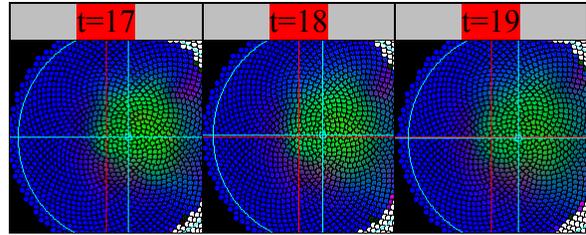
Illustration 115: Performance with Square

provides a memory lasting 11 ticks (from 8 thru 18).Similar examples can be found for the 2nd and 3rd level RFs.

4.4 General results and discussion

In this section we compare to alternative models, and discuss the behavioural emergence of memory-in-motion.

4.4.1 Square/Love and neurocentric models

As far as we know, all computational models that retain detailed visual information merely *hold* such information within retinotopic neurons and can demonstrate persistence over mostly static scenes. They do not actively *maintain* an internal model of the dynamic world, for example, by updating the spatial location of moving objects behind an occlusion or during an eye blink. Thus, we will show that they will have extensive problems with smear, especially under continuous motion. One such model is shown in Illustration 116.
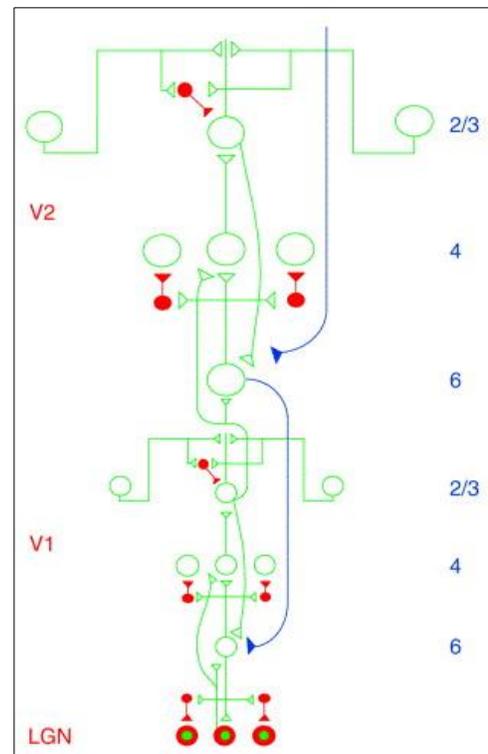
Illustration 116: LAMINART neurocentric model.
From (Grossberg, 2000), Fig. 4 showing bottom-up, top-down and lateral interactions within neural regions (red labels on left) and cortical layers (blue labels on right). Excitatory paths in green; inhibitory in solid red; attentional in blue.

If one assumes that each of the neurons above has a single state for integrate and fire, then converting it to a flow centric approach that facilitates maintenance and integration would require distributing such state to nearby neurons (splitting and joining of states) upon eye motion. There is no known way of doing so even with dendritic processing. The problem is that there are too many neural states oversimplified to lack spatial representations. ECM has no such problems as Emergic Units are stateless and stateful values represent their own spatial extent.

4.4.1.1 Lack of models

> *"Most models of vision focus either on the spatial or temporal aspects of visual processing and neglect the other component."* – (Rüter, Francis, Frehe, & Herzog, 2011)

The few models that claim to do both typically do not suffice. It is not enough for models to have internal dynamics, or temporal processing; they must also handle the external visual dynamics as well – they must handle self and object motion over space. Nevertheless, as far as we know, all such models have an evolutionary history that began with a static snapshot view of the world, and the handling of motion was a late afterthought. In this discussion, we will focus on a set of models originating for Stephen Grossberg as they demonstrate a long evolutionary history, have the most complex story, and have attempted to manage motion to some degree.
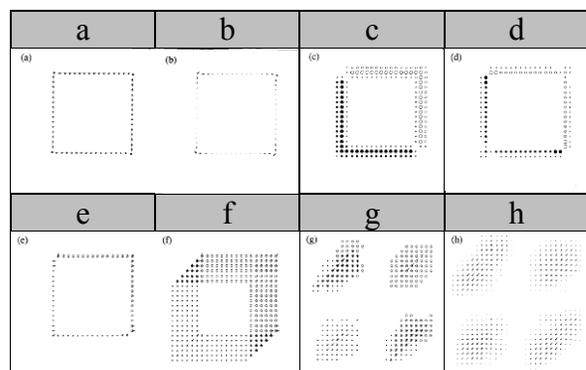
Illustration 117: Smear in multiple internal model locations (c,d,f,g,h) of FACADE
From (Grossberg & Mingolla, 1993), Fig. 31.

4.4.1.2   Smear in FACADE

In a FACADE based model of motion perception (Grossberg & Mingolla, 1993), a similar stimulus was used. A white square was moved diagonally over a grey background. The simulation results clearly demonstrate numerous internal points of smear as shown in Illustration 117.

In general, all of the Grossberg models make specific claims about how internal parts of their models interact to form conscious percepts of form and motion (e.g., Baloch & Grossberg, 1997). When there is no *actual* motion, but possibly *apparent* motion, the interpretation of results without smear is straightforward. However, without an actual realization of consciousness, it is difficult to interpret the smear inherent in their models. Clearly, under actual moderate motion, objects do not appear smeared to human participants. Having realized the problem of smear, an augmented model has been proposed (Francis & Grossberg, 1996). This will be further analyzed in the next section especially as it cannot eradicate the problem of smear under continuous motion.

There are two additional points to discuss. Firstly, this FACADE model has a scale parameter such that part (a) of Illustration 117 is one unit wide in terms of detecting motion. If motion was faster or slower, then this parameter would need to change. ECM has zero parameters. Furthermore, under a heterogeneous set of receptors where the size of the receptor varies with eccentricity, such a parameter is not even practical.

Secondly, ECM has no smear in this case because handling motion (currently only eye motion) is fundamental to its architecture and the information flow is motion invariant by design of its *maintain information coherence function*. ECM is not bound to neurons – it is *flowcentric*. However, FACADE is tied to specific retinocentric neurons at every point in its model. Where these have any forms of memories, smear must result under motion.

4.4.1.3 Boundary reset

The problem of smear was recognized by (Grossberg, 1991) and a computational solution was developed in (Francis et al., 1994; Francis & Grossberg, 1996; Grossberg & Grunewald, 2002). Simply, the old boundary information would be inhibited via a *reset* signal due to stimulus offset. As a consequence, new contour information could be created without
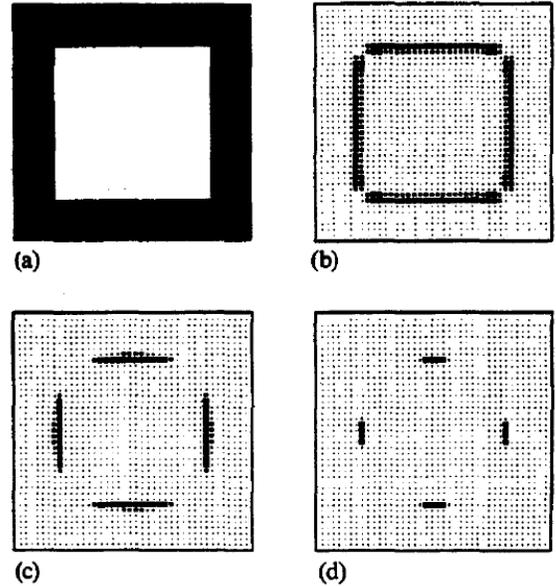
Illustration 118: Boundary erosion of BCS within FACADE
From (Francis et al., 1994), Fig. 5.

apparent smear. An example of the reset dynamics within the boundary-contour system (BCS) is shown in Illustration 118. This can be contrasted with the "boundary erosion" of ECM as shown in Illustration 119. Note that model phenomenology is similar between FACADE and ECM – the corners seem to be affected first. However, ECM is simpler in having neither an explicit boundary nor feature completion system, and certainly no reset mechanism.
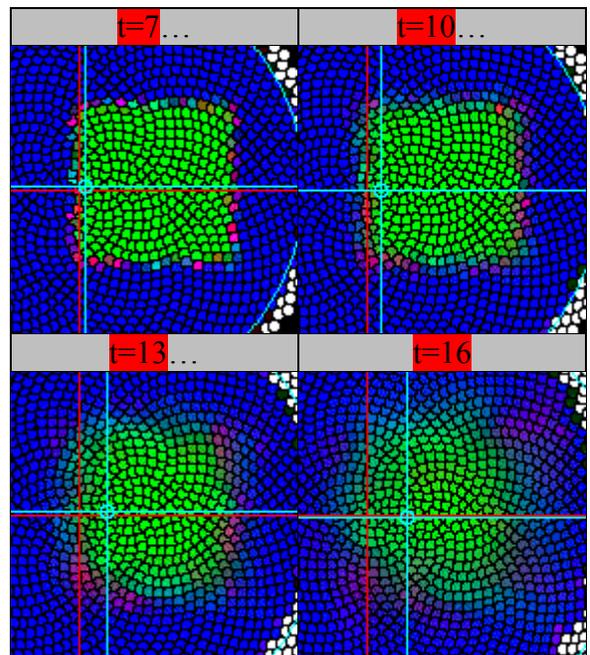
Illustration 119: Boundary erosion in ECM

Moreover, while ECM is explicitly a sensory only model that supports conscious percepts, in this particular case, the percepts are easier to interpret within ECM, and illustrate a problem of interpretation within FACADE. In FACADE and derivatives such as LAMINART (Grossberg, 2000; Raizada & Grossberg, 2003), percepts occur after the BCS and FCS modules have interacted. However, once boundaries start to erode, filling-in will dissipate as shown in Illustration 120. Thus, boundary erosion
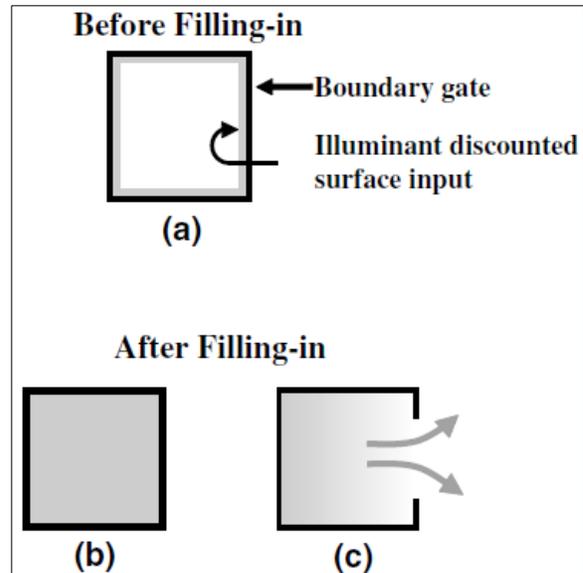
Illustration 120: Interpreting percepts with boundary gaps.
From (Grossberg & Yazdanbakhsh, 2005), Fig. 7. "Each boundary output to the surface system is accompanied by illuminant-discounted surface inputs which estimate the contrast magnitude across the corresponding edge. (a) Before filling-in. (b) If the boundary does not have a gap, it then can contain filling-in and may lead to a visible surface percept. (c) A boundary with gap lets the filling-in dissipate, thereby preventing a visible surface percept."

implies either a great smearing of colour (it is not clear in which direction, if not both), or a loss of visible surface perception. It is this final BCS/FCS interaction that needs to be demonstrated by simulation, and these are never shown for the Grossberg models in the case of actual motion. An example of the desired results without motion, but with flashing occlusions, is shown in Illustration 121.

In the case of continuous motion using our stimulus as an example, the leading edge of the green square over the blue background has the following characteristics – the blue RFs become 'black', i.e., they have a blue offset and a green onset. While the trailing edge has the opposite characteristics – the green RFs become 'black', i.e., they have a
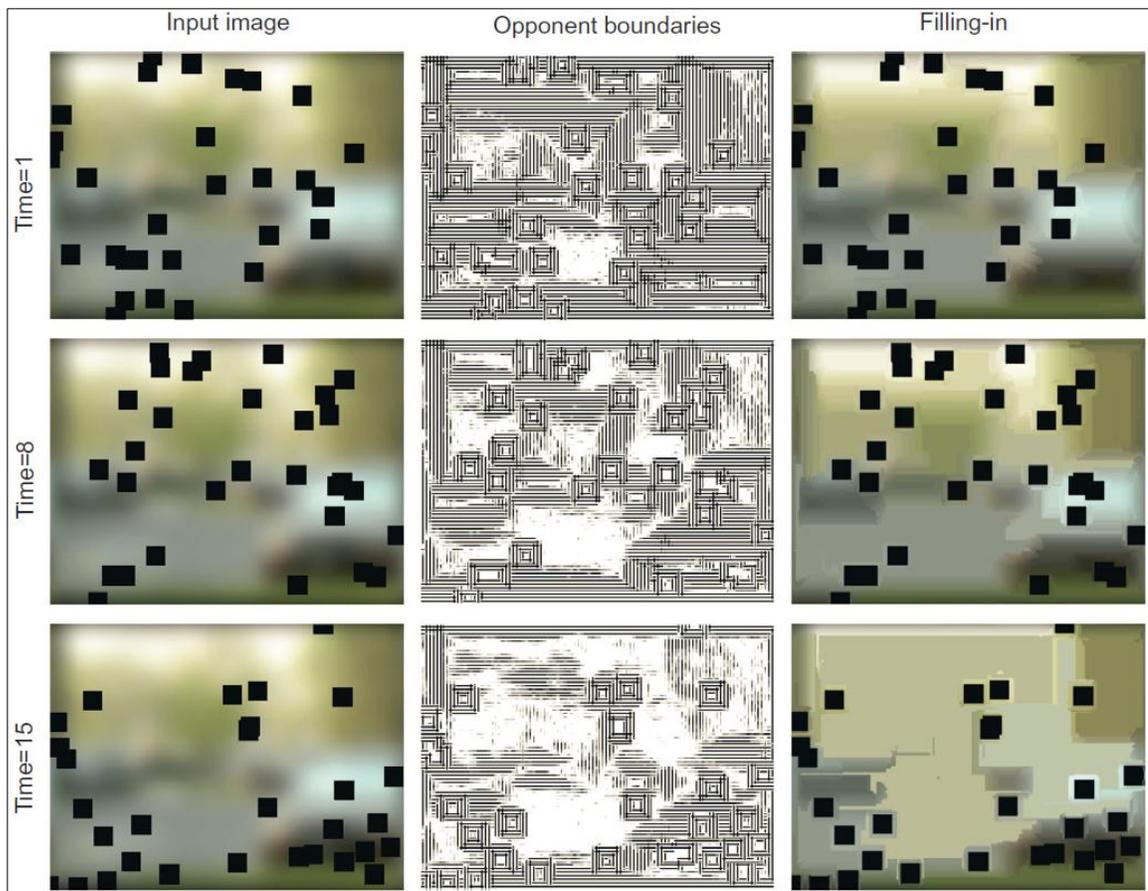
Illustration 121: Surface smearing after boundary erosion.
From (Francis & Kim, 2012), Fig. 4: Flashing squares over the input lead to faster boundary erosion that lead to greater surface smearing or "blur".

green offset and a blue onset. In terms of boundaries, these are continuously shifting.

This implies a continuous reset of the leading and trailing borders of the square, or

possibly a continuous oscillation. In the case of a continuous reset, this could further

imply that objects under continuous motion cannot be perceived as their surface features

continuously dissipate and are continuously filled-in by whatever background happens to

be at their leading and trailing edges. The case of oscillations is not phenomenological,

but might be a result of border competition. Regardless of outcome, if there are borders

then there is smear among borders, and if borders erode, then there is surface smear. The

Grossberg models cannot eliminate smear under continuous motion, and therefore do not

display such simulation results.

154

The purpose of reset signals is less to handle actual motion, and more to handle non-ecological stimuli that change from one long fixation period to another. It also highlights the sequential processing of perturbations within such models, despite the continuous unfolding of their internal dynamics. In FACADE, the BCS and FCS processing must occur first, before their interaction that lead to percepts. With stimuli changes across big periods of fixation, it makes sense to have concepts related to classes of perturbations such as image *onset* or *offset*. The old image is effectively erased and the new image is started afresh.

However, with a continuous *flowcentric* approach such as ECM, all information is used simply to maintain an internal model of the world taking motion and occlusions into account. There are no processing stages, only a single continuous flow. Capacity limitations and sequential processing are left to a future perceptual system. Information within neurons does not need to be reset as neurons do not hold information – the flow does. The flow never needs to be reset, because it is continuously maintained.

Therefore, the behaviour of the ECM and FACADE systems under motion is completely opposite. Under motion, FACADE exhibits smear, one way or another. ECM, on the other hand does not. Indeed, motion serves to undo the effects of dissipation and to refresh borders (p228). Under FACADE with continuous reset, it would be hard to explain the fact that hyperacuity remains the same with targets moving up to 2.5°/s or 300 photoreceptors/s (Westheimer & McKee, 1975).

We have no trouble with either motion or smear, nor have any reset behaviour. But then, we have no explicit memories nor have any explicit processes that deal with borders or surfaces.

The mechanism of reset to visual memory is identical to the mechanism of masking for iconic memory.

> *"In the natural visual environment, when we are making saccadic eye movements, each successive fixation is a potential mask for the retinal image of the previous fixation. Thus, if we had icons, they would successively mask each other. Such masking would help prevent successive overlapping fixations from obscuring each other by shortening the persistence and creating a blanking between each one (Breitmeyer 1980). This argument is useful in explaining the causes of saccadic suppression, but it does so at the expense of iconic storage. Maskability of persisting excitation reduces the potential usefulness of an icon by limiting its duration, perhaps to nothing." – (Haber, 1983)*

The concept of masking eliminates iconic memory, just as reset eliminates *neurocentric* models. Conversely, while masking is an actual effect, it is poorly understood in the context of continuous motion. The reset signal is completely understood as it is designed, but it cannot work under continuous motion.

4.4.1.4 Lack of smear in ECM

While the response at tick 16 (for the square stimulus) seems *blurred* or *smeared*, this is not so. Firstly, the

Illustration 122: Smearless "Flash memory" dissipates with Love

apparent smear was identical under the no-motion condition (not shown) and is therefore

due to dissipation. Secondly, as motion is steady, it can be seen that there is no smear for the first few ticks as shown in Illustration 122. Thirdly, and most importantly, smear is due to the piling-up of images during motion.

While there could be no piling-up after stimulus offset, these images can be contrasted with those of Test 1 of "3 Chapter:" with extracts shown in Illustration 123. It can clearly be seen that
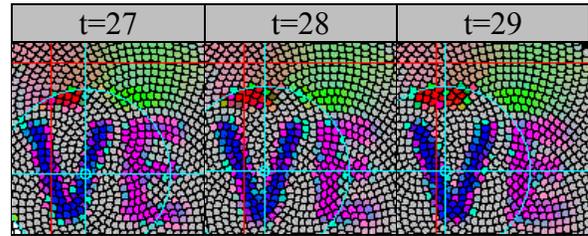
Illustration 123: Smearless transsaccadic memory dissipates with Love
Extract from Illustration 60 (p83)

while the letter 'V' moves within the photoreceptor region bounded by the cyan circular overlay (as well as part of the letter 'E'), then it does not *smear* – images of the letter 'V' at different points in space do not pile-up. However, when the letter 'E' moves beyond the reach of the photoreceptors, it begins to *dissipate*. Indeed, the apparent blurring of the letter 'E' is vertical while its motion was horizontal.

4.4.2   Flowcentric

Traditional models of visual processing are *neurocentric*, effectively storing detailed visual content within neurons (or reverberations between neurons). These neurons are invariably retinocentric. The
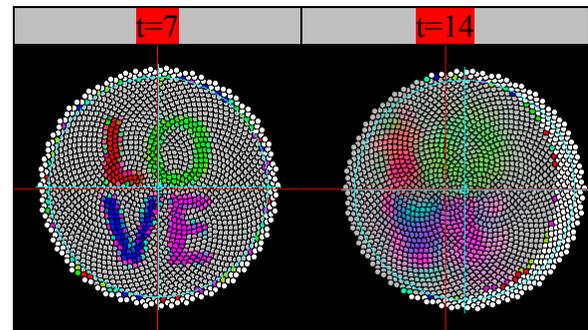
Illustration 124: Flowcentric Love

ECM architecture is *flowcentric*, effectively storing detailed visual content within flows of information. The flow of information has been engineered to be coherent, and because this includes the input flow from the environment, it is spatiotopic as well. These two

reference frames are shown in Illustration 124. The red axes indicate the spatiotopic reference frame. At tick 7, the red axes cross in the bottom-left corner of the letter 'O'. At tick 14, the red axes still cross in the bottom-left corner of the fading 'O'. However, the cyan axes represent the retinocentric reference frame. While they start identically, by tick 14, the gaze has shifted 7′ so that the cyan axes now cross at the bottom of the fading 'O'. The important aspect is that no specific part or memory of the letter 'O' is associated with any specific group of neurons. Instead, the memory of 'O' is retained in the flow of information.

Neurocentric and *flowcentric* models have different memory-like behaviours, especially under motion, and these differences are investigated in this and other chapters.

4.4.3   Differential flowcentric distortions

The spatiotemporal filling-in allows for a short (7 tick) memory by which time several different distortions have occurred. Colour is retained for all letters except the blue 'V' which seems to have been split into lilac, blue and cyan components. Fine detail has been lost which has affected letters 'L' and 'E'. The 'O' and 'V' retain their shape the longest. There is no consistent time constant to these distortions. The distortions may be affected by the spatial location within the sunflower mosaic pattern, but this has not been further investigated.

4.5   Conclusion

We have demonstrated the world's first cognitive model of memory-in-motion (or "flash memory") under moderate eye motion. This entails a spatiotemporal filling-in ability which is equivalent to maintaining a detailed internal visual model of the world. The

model is robust to stimulus type and requires no parameter changes. Indeed, the ECM architecture is parameter free.

The "flash memories" are conceptually stored in the flow of information that can remain within an RF level due to lateral connectivity. They are *flowcentric* and spatiotopic. They do not decay or fade away due to a time constant, but due to interference with other information in the flow. As different parts of the flow transit through a different path of heterogeneous neurons, they exhibit varying interference effects. Cohesion times do not have reliable time constants. This is a boon as a few emergic memories have the possibility to explain a plethora of memories distinguished by time factors.

Because the memories are *flowcentric*, they do not exhibit image smear and do not need their content to be erased. Visual models must take eye and object motion into account as a fundamental part of their architecture.

The phenomena of afterimages also implicate the maintenance of an internal model. For example, afterimages are stable under eye position with respect a reference point that appears to move, or vice-versa dependent on size of image (Davies, 1995; Pelz & Hayhoe, 1995).

4.6 Derivative eye blink model

The memory-in-motion model can be easily adapted to maintaining a stable view of the world during eye blinks and during moderate motion. While the analysis of visual stability concerns a different phenomenology with less talk of memory, the model is the

same. It can be found at Appendix C: An emergic model of filling-in during blinks on page 263 – actually, all our models are the same unified one.

4.7 Future

Currently we only handle eye motion. In the future, we intend to handle object motion and apparent motion. We also wish to validate our mechanism of dendritic processing and the possible role of its fine tuning via imagery under REM sleep.

Stimulus offset was simulated by effectively suppressing information from going upwards into the 1st level RFs. This was appropriate in relation to an implied layering mechanism that must be explicitly modeled. In particular, this mechanism would also need to account for the flash-lag effect. A stimulus that unexpectedly appears may be segmented into the occluding layer, before its motions is determined. A stimulus that unexpectedly disappears may remain in the fixational layer, while its occluder is immediately added to the occlusion segmentation layer.

5 Chapter: An emergic model of filling-in the blind spot

The eyes undergo incessant motion (even when fixating), yet most blind spot phenomenology is investigated with static eyes. While such phenomenology can be easily self-demonstrated (see Illustration 192 on page 232), most people are unaware of their own blind spot whether moving their eyes or not. It is not clear whether the models created for the static eye case under simple stimuli such as lines could handle the ecological case with dynamic eyes under natural stimuli.

In this chapter, we present a model of filling-in of the optical blind spot based on the Emergic Cognitive Model (ECM) architecture (p323). While ECM does not have any specific mechanism dedicated to filling-in, we hypothesized that both of its basic functions – *handling missing data*, and *maintaining information coherence* – could interact with experimental stimuli and measurements to cause this filling-in phenomena to emerge in a precise and timely manner. As part of characterizing the relative contributions of the two basic functions to the main phenomena, we also discovered that the model would cause border processing to emerge, and we analyze this effect in greater detail – it is a prerequisite for the precision of filling-in.

While ECM was designed with incessant eye-motion in mind, it must still handle the case of simple stimuli with a steady eye (jitter cannot be stopped).

5.1 Introduction

Much of the literature surrounding the optical blind spot is reviewed in Appendix B: *An emergic model of filling-in the foveal blue scotoma* on page 231. It will not be repeated here except for the summarizing the following points. Firstly, while there are similarities between the two scotomas, the foveal blue scotoma is only missing S (blue) cones and could interpolate from nearby L (red) and M (green) cones. The optical blind spot is missing all photoreceptors (even rods) and must fill-in spatiotemporally over distances much larger than one cone. Secondly, there are no known models for the blue scotoma, but several for filling-in over the blind spot.

Thirdly, the neurobiology behind the blue scotoma is unknown as the region is quite small and uncharted. However, much more is known about the blind spot. There are receptive fields in V1 behind the blind spot that can mapped from the contralateral eye (Komatsu, Kinoshita, & Murakami, 2002; Komatsu, 2011), from stimulus near the blind spot (Komatsu, Kinoshita, & Murakami, 2000), or can be interpolated by sweeping bars longer than the blind spot across it at various

Illustration 125: Large RFs within blind spot
Broken line represents blind-spot, while ellipses represent RFs determined (Komatsu et al., 2002).

orientations (Forani Jr. et al., 2003). Some of these RFs are extremely huge as shown in Illustration 125. This may explain advantages such as better spatial (Vernier) alignment over the blind spot than on the equally eccentric mirror side of the retina (Crossland & Bex, 2009). Their huge size may be due to massive feedback projections from V1 to the

LGN that have broader arborization than feedforward projections from the LGN to V1 and this looping could be one mechanism for blind spot filling-in (Yokoi & Komatsu, 2010). We, of course, are greatly interested in all forms of recurrence including attentional matters. For example, blind spot filling-in is easier when covert attention around the blind spot is diffused rather than focused (Lou & Chen, 2003). Indeed many related phenomena imply recurrence from higher regions. For example, periodic stimuli are more effective, and even patterns can be filled-in (R. J. Brown & Thurmond, 1993; Kawabata, 1982) including appearance of illusory objects (Araragi, Ito, & Sunaga, 2008); features seem to be grouped based on similarity then segregated before completion (Kawabata, 1984); the motion aftereffect can be transferred from one eye into the blind spot of the other (Murakami, 1995); and the motion of a moving object can be extrapolated into the blind spot (Maus & Nijhawan, 2008). While knowledge of blind spot neurobiology is easy to find, the mechanisms behind the underlying surface filling-in are poorly understood (Komatsu, 2011)

> *"there is no agreement even on the basic question whether neural activity across the retinotopic representation of surfaces correlates with the surface percept. ...divergent results [snip] are compounded with blurred theoretical concepts, and a virtual absence of direct evidence on the mechanisms that might underlie the perception of surfaces."* – (De Weerd, 2011)

Finally, general models of filling-in or completion *beyond* the blind spot (e.g., Cohen & Grossberg, 1984; Kalar et al., 2010) cannot suffice *within* the blind spot as the phenomenology differs between the two regions. For example, invisible contours are not completed through the blind spot (Maertens & Pollmann, 2007); surfaces that are filled-in through the blind spot, must cover at least two quadrants (Kawabata, 1982); and visible

contours may be only partially completed if they do not extend enough beyond the blind spot (Abadi, Jeffery, & Murphy, 2011).

5.2 Emergic Model

In order to explain the phenomena of this chapter, we have taken the ECM architecture as summarized on page 5, and made no changes. For the filling-in of the optical blind spot, we have two intuitions as to how ECM might explain the results.

For static eyes, ECM has a slow filling-in behaviour that appears like diffusion, but respects borders and is best characterized as interpolative (between borders) rather than extrapolative (beyond border) (p57). Larger sized RFs improve diffusion times, but in these experiments, we verify that larger RFs can also help to fill-in over the extended regions of space belonging to the scotoma – can larger RFs convert extrapolation into interpolation?

For dynamic eyes, ECM has an instantaneous filling-in behaviour due to the shifting of information from one region of space into another due to incessant eye motion (p171). In these experiments, we verify that a robust set of stimuli can be shifted into the optical blind spot.

The *handling missing data* function would be invoked in both cases. When the eye is stationary, top-down interpolated values will be used. When the eye is moving, lateral shifted values will be used. The *maintaining information coherence* function performs the shifting of information to compensate for eye movement. It also requires values to have

positional information, and it is this that enables the emergence of border effects. In these experiments, we verify whether borders are indeed respected when filling-in.

5.3 Test 1-3: Dynamic blind spot over robust stimuli

While not all stimuli are filled-in within the blind spot under static eye conditions (Kawabata, 1982), it appears as if all stimuli are filled-in under dynamic eyes. In this test we characterize the memory-in-motion behaviour of the ECM architecture as the blind spot moves through the salient aspects of an image. To demonstrate robustness of this spatiotemporal filling-in phenomenon, the model was examined across a wide range of stimuli.

5.3.1 Methods

The methods used are identical to those of page 80 except as indicated herein.

5.3.1.1 Subject analysis summary

A similar virtual agent as in Test 1-4 of "4 Chapter:" (p143) was used, except that there was no blinking, and an optical blind spot was added. This is summarized in Illustration 126 while details can be found in the

Illustration 126: Photoreceptor mosaic of agent used in blind spot tests 1-3

supplement (s163). The visual angle for the entire photoreceptor mosaic was 88.51′ while the visual angle for the 1st, 2nd and 3rd RF levels was 98.86′, 104.18′ and 110.59′ respectively.

5.3.1.2   Stimuli

Similar stimuli as in tests 1-3 of "3   Chapter:" (p81) were used and are shown as viewed by *person* in Illustration 127, Illustration 128 and Illustration 129. They are fully described in the supplementary material (s166, s167).
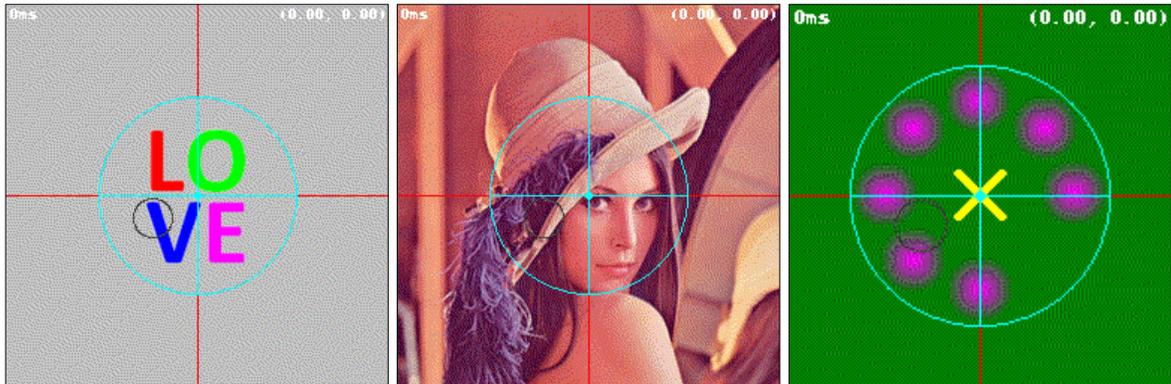
Illustration 127: Blind spot Test 1 stimulus (Love) with overlay   **Illustration 128: Blind spot Test 2 stimulus (Lena) with overlay**   **Illustration 129: Blind spot Test 3 stimulus (Chaser) with overlay**

5.3.1.3   Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to move its eye in the following pattern that would cause the blind spot to traverse the most salient aspects of the image.

- Move right for 8 ticks at a speed of $4'$/tick ($\pm$ jitter)

- Move up for 8 ticks at a speed of $4'$/tick ($\pm$ jitter)

- Move left for 8 ticks at a speed of $4'$/tick ($\pm$ jitter)

- Move down for 8 ticks at a speed of $4'$/tick ($\pm$ jitter)

The artificial brain was monitored for the first 32 time ticks. Note that there is an eye movement planning and distribution delay of 6 ticks before the eye actually moves as desired, so the above plan will not complete.

For Test 2, the eye moved in a slightly different path to ensure the blind spot traversed salient parts of the stimulus (eyes, nose and mouth).

5.3.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s172, s183, s194). When presented as an animation (Leibovitz, 2012a), all the results for this chapter can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=377
```

In this section, we summarize and discuss various highlights.

5.3.2.1   Blind spot filling-in behaviour

The dynamic behaviour of filling-in the blind spot at the 1st level RFs are summarized in Illustration 130, Illustration 131 and Illustration 132. Eye jitter starts at tick 7 and is indicated via a grey title background.
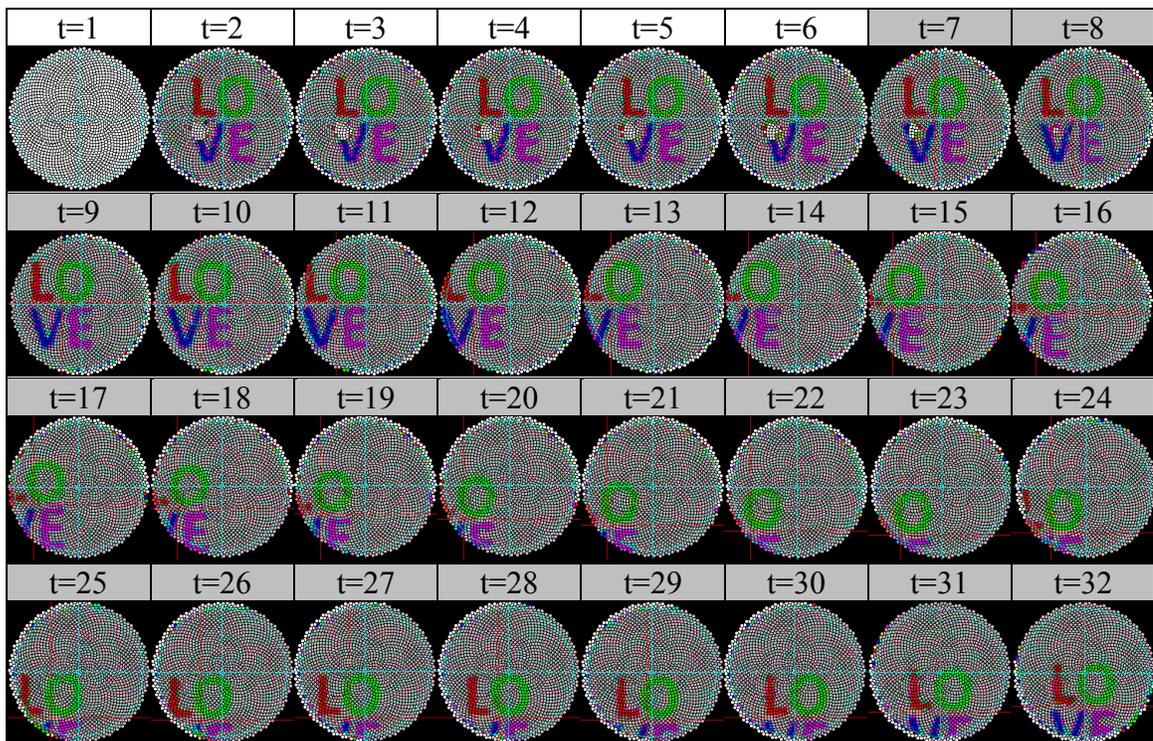
Illustration 130: Blind spot filling-in behaviour for Test 1 (Love)
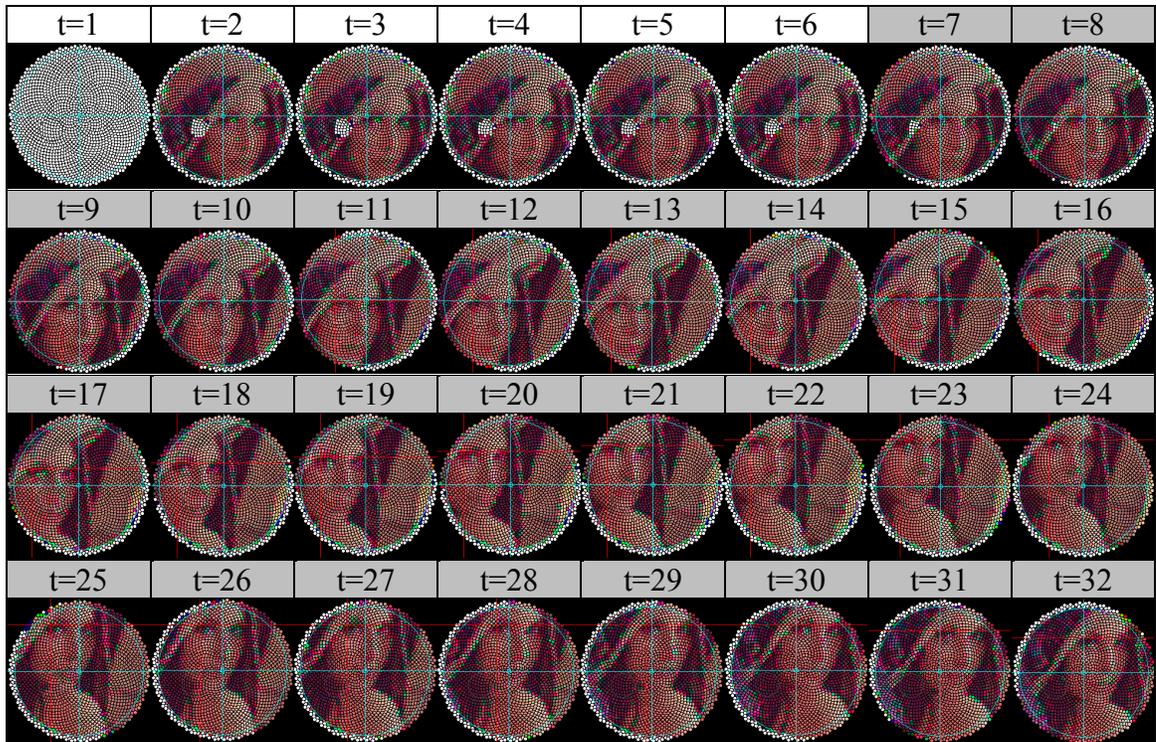
Illustration 131: Blind spot filling-in behaviour for Test 2 (Lena)
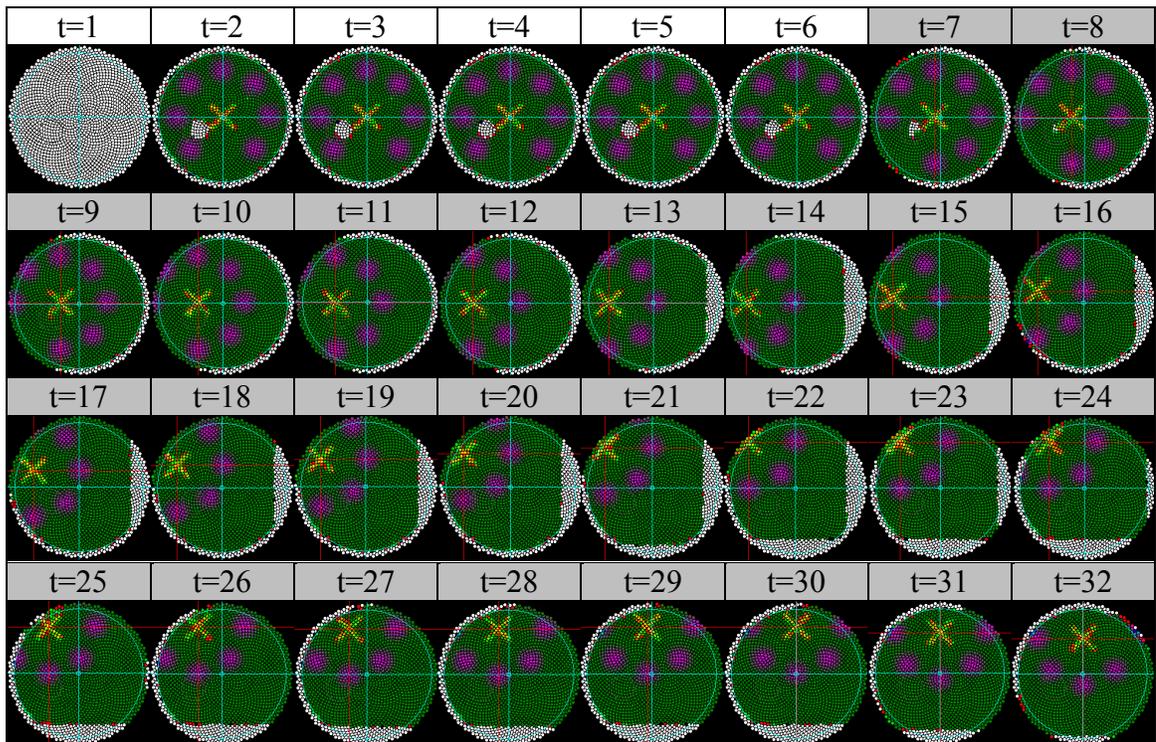
Illustration 132: Blind spot filling-in behaviour for Test 3 (Chaser)

Similar spatiotemporal filling-in behaviour occurs at the 2^{nd} and 3^{rd} RF levels. These details are available in the supplemental material.

168

After tick 8, it is extremely difficult to find artefacts due to the blind spot. It has effectively disappeared under incessant eye motion even when traversing salient parts of the stimulus.

5.3.2.2 Static blind spot filling-in performance

The first six ticks do not include eye motion and can demonstrate blind spot filling-in behaviour under static eye conditions. These are shown for the 1^{st} Level RFs in Illustration 133, Illustration 134 and Illustration 135.
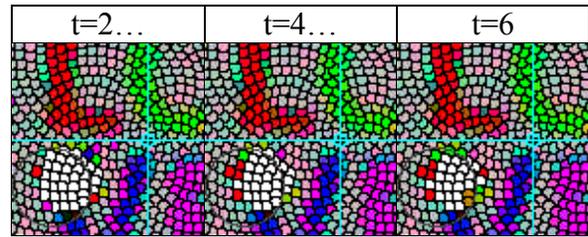
Illustration 133: Static blind spot filling-in performance for Test 1 (Love)

Similar spatiotemporal filling-in behaviour occurs at the 2^{nd} and 3^{rd} RF levels. These details are available in the supplemental material.
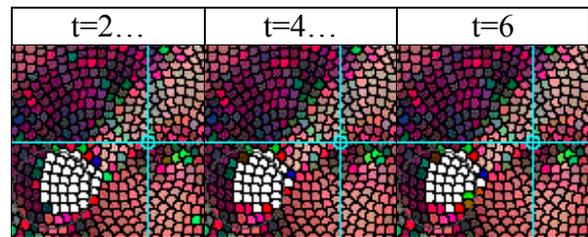
Illustration 134: Static blind spot filling-in performance for Test 2 (Lena)

In these configurations it does not appear as if static blind spot filling-in will ever complete demonstrating that ECM does not have an explicit diffusion process.
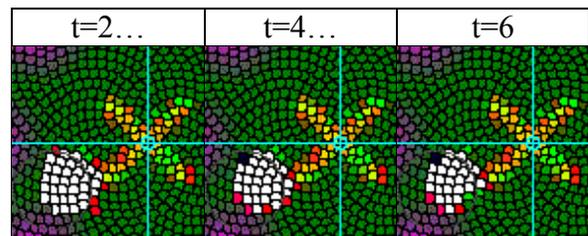
Illustration 135: Static blind spot filling-in performance for Test 3 (Chaser)

5.3.2.3   Dynamic blind spot filling-in performance

The eyes start to move at tick 7 and this can be used to demonstrate blind spot filling-in behaviour under dynamic eye conditions. These are shown for the 1st Level RFs in Illustration 136, Illustration 137 and Illustration 138.
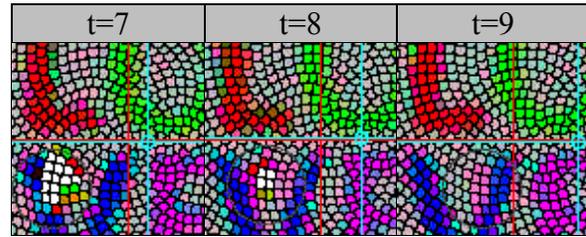
Illustration 136: Dynamic blind spot filling-in behaviour of 1st Level RFs with Love stimulus

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.
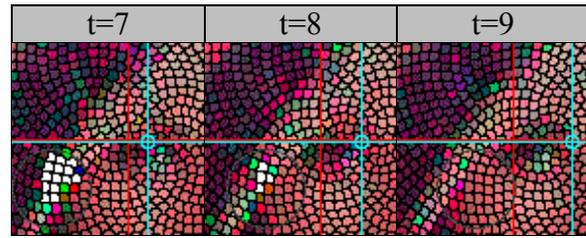
Illustration 137: Dynamic blind spot filling-in behaviour of 1st Level RFs with Lena stimulus

In these configuration, it took 2 ticks to fill-in the blind spot completely but that could be due to eye movement that was only ½ the size of the blind spot. Instantaneous filling-in will be demonstrated next.
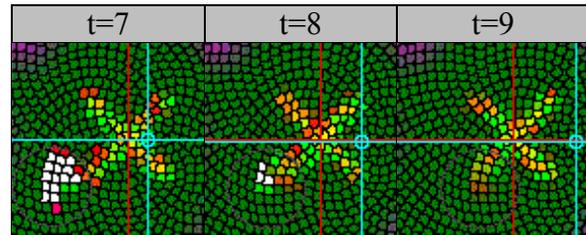
Illustration 138: Dynamic blind spot filling-in behaviour of 1st Level RFs with Chaser stimulus

After tick 8, it is extremely difficult to find artefacts due to the blind spot. It has effectively disappeared under incessant eye motion even when traversing salient parts of the stimulus.

5.3.2.4 Instantaneous dynamic filling-in performance

Once the blind spot has been filled-in, then this can be used to demonstrate that subsequent filling-in of static stimuli due to eye motion is instantaneous. This is detailed for the 1st Level RFs in Illustration 139, Illustration 140 and Illustration 141. Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

For Test 1 at tick 26, the blind spot is moving leftwards but is not yet over the letter 'L'. This is shown by the top row of the illustration showing the photoreceptor activity. At tick 27, part of the blind spot has moved over the letter 'L'. At this point, the *missing* photoreceptors would
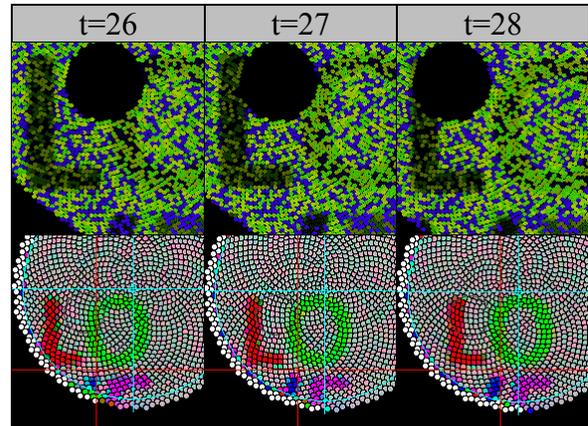
Illustration 139: Instantaneous dynamic blind spot filling-in behaviour of 1st Level RFs

have sent the red information pertaining to the letter 'L' up to the 1st level RFs, although this would incur a 1 tick delay. However, they do not exist so no information was sent upwards. Nevertheless, the 1st level RFs already have the correct red information pertaining to the letter 'L' within the blind spot.

For Test 2 at tick 11, 15 and 19, the blind spot was over the eye, nose and mouth respectively. If it even took a single tick delay, then the blind spot filling-in would be out of alignment with the rest of the scene due to incessant eye motion. However, the blind spot is filled-in instantaneously with no parts out of alignment.
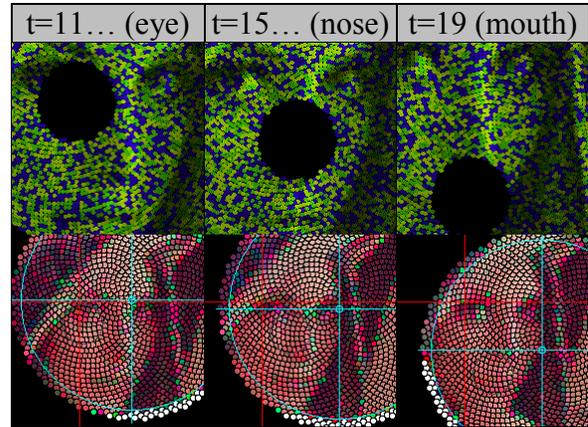
Illustration 140: Instantaneous dynamic blind spot filling-in behaviour of 1st Level RFs

For Test 3 at tick 10, 16 and 25, the blind spot was over the fixation point, a lilac and another lilac respectively. If it even took a single tick delay, then the blind spot filling-in would be out of alignment with the rest of the scene due to incessant eye motion. However, the blind spot is filled-

Illustration 141: Instantaneous dynamic blind spot filling-in behaviour of 1st Level RFs

in instantaneously with no parts out of alignment. Indeed, with stimuli having gradient edges, it is even harder to detect any visual artefacts.

5.3.3   Discussion

We have demonstrated a cognitive model of filling-in the optical blind spot with dynamic eyes. The model is robust to natural stimuli. During static eye conditions, filling-in does not complete indicating that ECM does not have explicit diffusion or a filling-in

processes. We hypothesize that having larger 2^{nd} and 3^{rd} level RFs will facilitate static

filling-in, and this will be tested in experiments 4-7.

The emergic filling-in with dynamic eyes is instantaneous, incurring no time delay. This

is because the information has been *predictively* shifted into the blind spot from other

lateral RFs. Such predictive shifting overcomes processing delays, and is a recognized

function of prediction (Lim & Choe, 2008; Nijhawan, 2002). Nevertheless, *prediction* is

possibly not the best term in these circumstances because the system knows ahead of time

how the eye would move – eye movement is *planned* in advance, not predicted. Indeed,

the very reason for planning motor movements in advance is to enable instantaneous

coherentizing of information. We prefer to use the term *prediction* to extrapolate a future

observable from previous measures. In this situation, there have been no measurements

and no extrapolation function. *Maintaining information coherence* is simply our

functional name for the observed behaviour.

However, in Illustration 136 it appears that filling-in was not instantaneous taking 2 ticks.

We argue that if the eye shifted by an amount greater than the size of the blind spot, and

if lateral connectivity allowed, than this filling-in would indeed have been instantaneous.

Conversely, if eye motion was $1/10^{th}$ the size of the blind spot, then it would take 10 ticks

to fill in. While such a procedure was not specifically tested, instantaneous filling-in was

demonstrated in Illustration 139. We also surmise that if eye motion was greater than

lateral fan-out connectivity, then the information would be effectively masked, and the

blind spot would reappear. Such masking was demonstrated in Illustration 75 (p96) .

Under ordinary circumstances, once the blind spot has been filled-in, eye-motion ensures that it remains filled-in accurately even when moving across salient information such as edges. It is no wonder that we are ordinarily unaware of our optical blind spot. The disappearance of the blind spot under motion, and the extreme difficulty in detecting its artefacts is much more apparent in the animations (Leibovitz, 2012a) available at http://emergic.upwize.com/?page_id=377. Careful analysis of image frames one-by-one can indeed find these artefacts.

5.4   Test 4: Static eyes over static line

In Tests 1-3, it was found that the filling-in of the blind spot was accurate and instantaneous under dynamic eye conditions. However, under traditional static eye conditions, it did not complete. We hypothesized that having larger 2^{nd} and 3^{rd} level RFs will facilitate static filling-in. In the next three tests we validate our hypothesis using traditional blind spot stimuli. In this test, the static blind spot will be situated directly over a static line (or bar) stimulus, as in column (d) of Illustration 153 on page 184. This is the 1^{st} of three traditional blind spot experiments used to test our model.

5.4.1   Methods

The methods used are identical to Test 1 except as detailed herein.

5.4.1.1 Subject analysis summary

A virtual agent was constructed as shown in Illustration 142. Its blind spot was centered where control of higher level RF sizes is greatest. It is fully described in the supplementary material (s164). The visual angle for the entire photoreceptor mosaic was 69.96′ while the visual angle for the 1^{st}, 2^{nd} and 3^{rd} RF levels was 100.92′, 143.01′ and 113.34′ respectively.
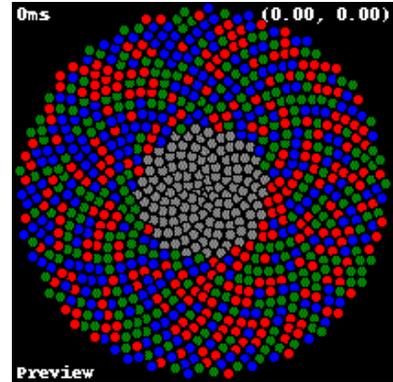
Illustration 142: Photoreceptor mosaic of agent used in blind spot tests 4-7

5.4.1.2 Stimuli

Traditional tests use a black bar over a white background (Ramachandran, 1992a). We chose a red coloured bar over a dark background as shown in Illustration 143. It is fully described in the supplementary material (s168).
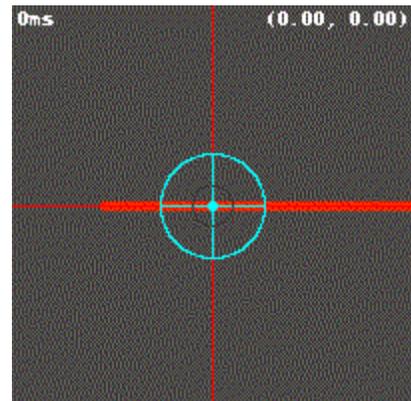
Illustration 143: Blind spot Test 4 stimulus (LineStatic) with overlay

5.4.1.3 Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to keep its eye steady (barring involuntary jitter). The artificial brain was monitored for the first 32 time ticks. Note that there is an eye movement planning and distribution delay of 6 ticks before the eye actually starts to jitter.

5.4.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed
in the supplementary material (s205). In this section, we summarize and discuss various
highlights.

5.4.2.1 Blind spot filling-in behaviour

The model's behaviour of blind spot filling-in for the 1st Level RFs is summarized in
Illustration 144. Eye jitter starts at tick 7 and is indicated via a grey title background.
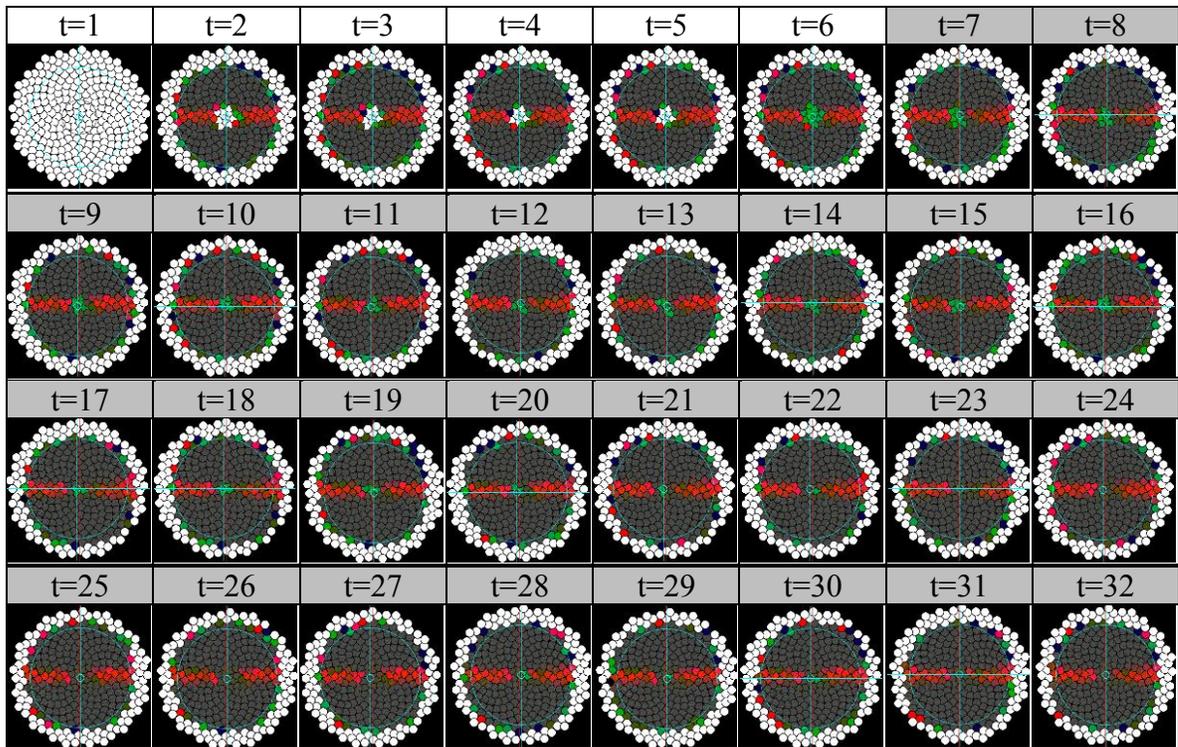
Illustration 144: Blind spot filling-in behaviour of 1st Level RFs with Static Line stimulus

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These
details are available in the supplemental material.

5.4.2.2 Filling-in performance details

Details of filling-in the 1st Level RFs (Illustration 145) produce interesting model artefacts. The first colour to be filled-in is green rather than red. This is surprising as there is no green in the

t=5	t=6…	t=24
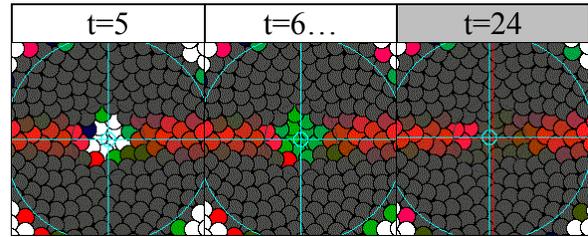
Illustration 145: Static blind spot filling-in details of 1st Level RFs

stimulus. By tick 24, the green is eradicated and there is a hint of redness that is filled-in, but this is far from "instantaneous". At best, this could be thought of as partial filling-in. Of great import, whatever has been filled-in respects the horizontal boundaries of the stimulus.

Details of the 2nd Level RFs (Illustration 146) produce better performance results. The hint of redness is available by the first stimulus volley (t=3) and the center completes by the next tick (t=4).

t=2	t=3	t=4
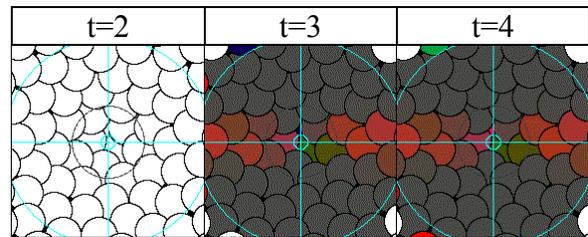
Illustration 146: Static blind spot filling-in details of 2nd Level RFs

Details of the 3rd Level RFs (Illustration 147) produce best performance results. The hint of redness is available by the first stimulus volley (t=4) and it is already complete.

t=3	t=4	t=5
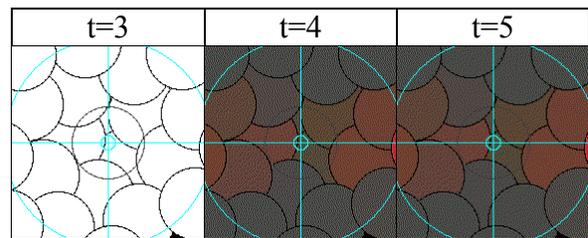
Illustration 147: Static blind spot filling-in details of 3rd Level RFs

177

5.4.3   Discussion

5.4.3.1   Green filling-in?

In Illustration 145, green was initially filled-in rather than red, but it was eventually eradicated. This is a stochastic artefact due the heterogeneity of cone LMS colour sensitivity types (p15). Indeed, in looking at all model images throughout this thesis, one will find random colours on the borders of the data-less RFs shown in the GUI as white. For the RF at tick 5 in Illustration 145, there appears to be 3 green-only patches, 1 blue only patch, and 1 red only patch that have partially infiltrated the blind spot. We assume that green won this initial lateral filling-in, but it was eventually eradicated by top-down filling-in. The weights of lateral and top-down influence have not been optimized, but it still takes more time for information to flow up and then back down. We have not analyzed this behaviour any further.

5.4.3.2   Partial filling-in

There is a hint of red that has fully filled-in across the blind spot (t=24 in Illustration 145). But currently, the model can charitably be described as only partially filling-in. Is the model incorrect? Can something be added to it so that filling-in completes? We will analyze some options next.

Firstly, partial filling-in is a recognized blind spot phenomenon when the bars do not extend far enough away from the edges of the blind spot (Abadi et al., 2011). Models that *always* fill-in are just as incorrect as models that *always* only partially fill-in. The question is where such contextual information is processed? There are numerous blind

spot phenomena implicating recurrence from higher levels of processing. For example, periodic stimuli are more effective, and even patterns can be filled-in (R. J. Brown & Thurmond, 1993; Kawabata, 1982) including appearance of illusory objects (Araragi et al., 2008); features seem to be grouped based on similarity then segregated before completion (Kawabata, 1984); the motion aftereffect can be transferred from one eye into the blind spot of the other (Murakami, 1995); and the motion of a moving object can be extrapolated into the blind spot (Maus & Nijhawan, 2008). So perhaps our sensory model is adequate, and it is higher level sensory or perceptual processing that completes the filling-in. We will gather more data in Tests 5 and 6 coming up with one solution in Test 7.

Secondly, it is important to note that our RFs are not visual representations, but hold statistical information for each LMS colour independently. It is only the GUI that combines this information into a uniform colour. For example, using the more familiar RGB notation, with r being the radius of the blind spot, the dark background could be thought of as B=5% @(x=±4r, y=±4r), green could be unspecified, and the horizontal line would be R=100%@(x=±4r, y=±2′). However, as darkness covers a greater area, even though the statistics for the red line are retained, the GUI would show the entire RF as dark with a tinge of redness. Thus, the GUI is a poor measure of underlying information flows in some cases. However, it is due to these statistical properties, that howsoever much red is filled-in, it respects borders.

Thirdly, ECM is a sensory model. While the 1st level RFs do not complete the filling-in of the red horizontal line, they are completed in the 2nd and 3rd RF levels (Illustration 146 & Illustration 147). Who is to say how conscious perceptual processing combines all such information flows? Indeed, if we did show a properly filled-in level with isomorphic neural counterparts, then that would amount to an invalid Cartesian Theatre.

Fourthly, while the larger sized RFs did facilitate filling-in, perhaps they too were not optimal. They could be the wrong size, but more importantly, they could be poorly placed. The 3rd layer contains only 22 RFs which can be seen in Illustration 147. However, the innermost most five are shown in Illustration 148. Note that none of the RFs is centered under the blind spot indicated with a
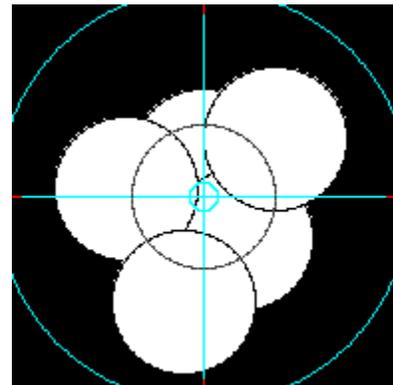
Illustration 148: Five innermost RFs in 3rd level

black circle as an overlay (it's hard to tell it is not an RF). The problem with not having a centred large RF is that ECM can interpolate between points (which could allow for filling-in if centered), but cannot extrapolate (p57). This particular placement optimization is not attempted, as an alternative solution will be demonstrated in Test 7.

Finally, it is technically ambiguous what colour ought to be filled-in. Just as there is a red stimulus on both side of the blind spot, so too there is a dark stimulus on opposite sides, indeed a lot more. If two long bars form a cross with an intersection at the blind spot, only the longer bar will be completed (Ramachandran, 1992a). This again indicates

global contextual processing. There are numerous other factors that may affect such ambiguities (R. J. Brown & Thurmond, 1993).

5.4.3.3   Alternative model

Satoh & Usui (2008a, 2008b, 2009) present a mathematical model for filling-in the blind spot based on this specific stimulus, but it requires a priori knowledge about what constitutes missing information. Then they apply their algorithm for digital image inpainting purposes with non-biological results as shown in Illustration 13 on page 22.

5.5   Test 5: Static eyes over moving line

This is the 2nd of three traditional blind spot experiments used to test our model. One way to map the V1 RFs over the blind spot is to sweep bars longer than the blind spot across it. The bars can be at various orientations, and where these intersect, that indicates the spatial extent of
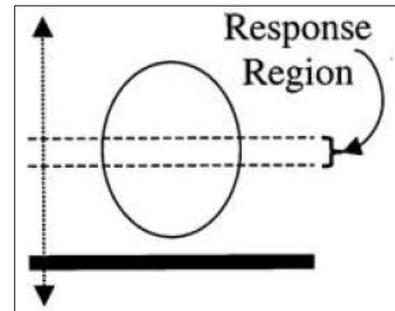
Illustration 149: Moving a large bar across the blind spot
From (Forani Jr. et al., 2003).

the RF. A sweep of a horizontal bar is shown in Illustration 149. In this test, we reproduce such a phenomenon within the model.

5.5.1   Methods

The methods used are identical to Test 4 except as detailed herein.

5.5.1.1   Stimuli

The stimuli is similar to that of Test 4 except that it was positioned slightly above the blind spot as shown   in Illustration 150 so that it could be moved downwards across the blind spot by 1 pixel per tick. It is fully
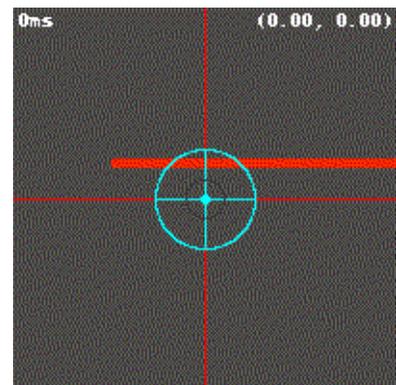
Illustration 150: Blind spot Test 5 stimulus (LineAcross) with overlay

181

described in the supplementary material (s168).

5.5.1.2   Procedure

While the agent procedure was identical to Test 4 (agent kept its eye static), the

horizontal bar was moved downwards by 1 pixel per tick.

5.5.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed

in the supplementary material (s216). In this section, we summarize and discuss various

highlights.

5.5.2.1   Blind spot filling-in behaviour

The model's behaviour of blind spot filling-in for the 1^{st} Level RFs is summarized in

Illustration 151.
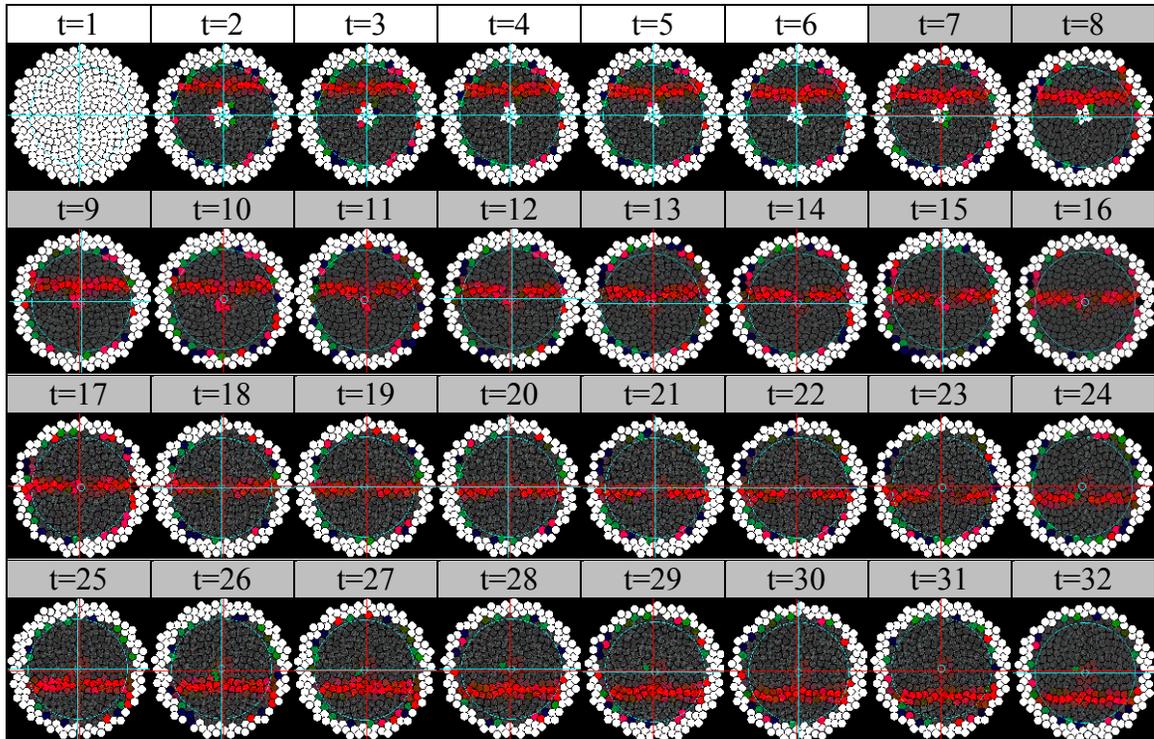
Illustration 151: Blind spot filling-in behaviour of 1^{st} Level RFs with Moving Line stimulus

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

5.5.2.2   Filling-in performance

Details of filling-in the 1st Level RFs (Illustration 152) produce interesting model artefacts. At tick 9, red appears to predictively fill-in the blind spot.
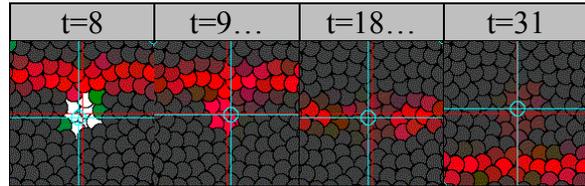
Illustration 152: Static blind spot filling-in performance

At tick 18, the horizontal bar has completed within the blind spot, however it is still faint, but better than Test 4. At tick 31, there is a faint memory trace of the redness left behind. It appears that the dark background cannot fill-in either by tick 8 or 31.

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

5.5.3   Discussion

The predictive red filling-in at tick 9 of Illustration 152 can be accounted for by ECM in the same manner as green filling-in of Test 4 – as due to photoreceptor heterogeneities at the boundary of the blind spot and relation to RF sizes and placement. While ECM can compensate for self-motion, at this time it does not track object motion, nor predict object locations. The term *prediction* while descriptive has no basis in the explicit computations performed – of compensating for known planned eye motion. Perhaps it is an undesired artefact. Further analysis of this *predictive* effect (or artefact) has not been performed.

By tick 18, the horizontal bar is across the middle of the blind spot and appears to be filled-in correctly. While faint, it is more salient than in Test 4. We have taken little effort to optimize the size and placement of RFs. Taken together, this constitutes modest support to our hypothesis that larger RFs facilitate filling-in. We will await Test 7 for a better solution.

By tick 31, there is still a faint memory of redness at the blind spot. This comes with some irony. While ECM has no smear due to memory-in-motion under dynamic eye conditions (p156), it has object smear for objects in motion under static eye conditions as shown here. Certainly, in the future we plan on handling object motion. It could also be that intermittent microsaccades could also wipe away such memories due to memory-in-motion. There is certainly an interference effect at the blind spot with multiple stimuli (red or dark) competing to fill-in. In Test 7, this interference is handled.

5.6 Test 6: Static eyes over inserted line

This is the 3rd of three traditional blind spot experiments used to test our model. Matsumoto & Komatsu (2005) have shown that as a bar stimulus is inserted into the blind spot, it is eliminated from perception until the bar crosses to the other side of the blind spot, at which time the whole bar is filled in. This is shown in Illustration 153. This would be an
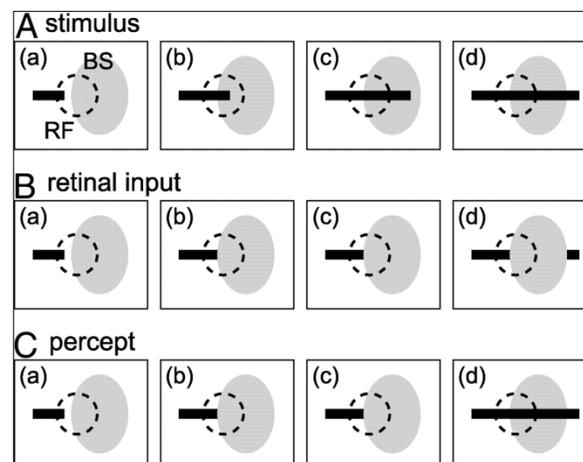
Illustration 153: Bar into blind spot
A bar stimulus is inserted into the blind spot. Each column represents a different length of insertion. Row A show the stimulus; row B the retinal input; while row C is the percept. From (Matsumoto & Komatsu, 2005), Fig. 1.

example where object prediction does not occur.

5.6.1   Methods

The methods used are identical to Test 4 except as detailed herein.

5.6.1.1   Stimuli

The stimuli is similar to that of Test 4 except that it was

positioned slightly to the right of the blind spot as shown

in Illustration 154 so that it could be inserted leftwards

into the blind spot by 1 pixel per tick. It is fully described

in the supplementary material (s169).
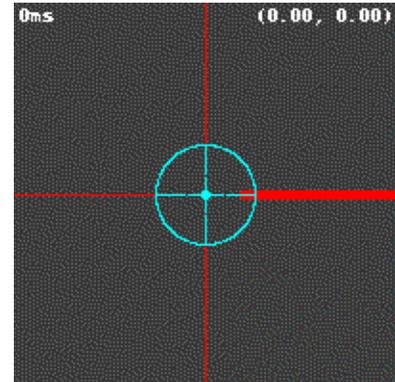
Illustration 154: Blind spot Test 6 stimulus (LineIn) with overlay

5.6.1.2   Procedure

While the agent procedure was identical to Test 4 (kept its eye static), the horizontal bar

was moved leftwards by 1 pixel per tick.

5.6.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed

in the supplementary material (s227). In this section, we summarize and discuss various

highlights.

5.6.2.1   Blind spot filling-in behaviour

The model's behaviour of blind spot filling-in for the 1st Level RFs is summarized in

Illustration 155.

t=1	t=2	t=3	t=4	t=5	t=6	t=7	t=8

t=9	t=10	t=11	t=12	t=13	t=14	t=15	t=16

t=17	t=18	t=19	t=20	t=21	t=22	t=23	t=24

t=25	t=26	t=27	t=28	t=29	t=30	t=31	t=32
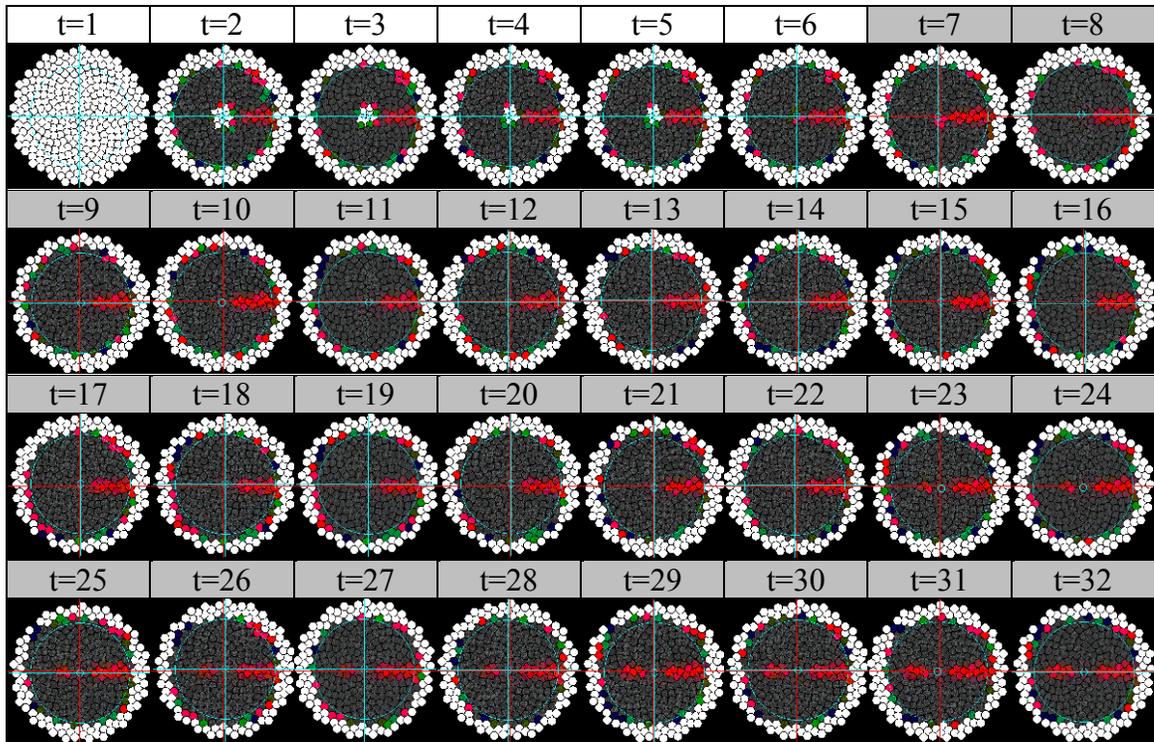
Illustration 155: Blind spot filling-in behaviour of 1st Level RFs with Inserted Line stimulus

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

5.6.2.2   Filling-in performance details

Details of filling-in the 1st Level RFs is shown in Illustration 156. Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

t=22	t=23…	t=32
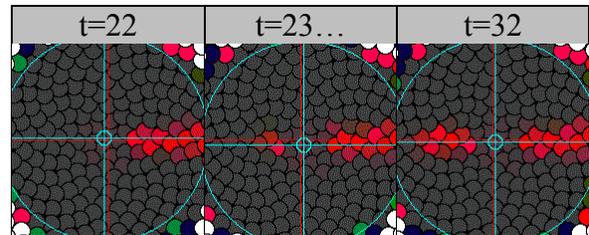
Illustration 156: Static blind spot filling-in details of 1st Level RFs

Up to tick 22, the bar is being swallowed by the blind spot. At tick 23 it appears on the other side, and by tick 32 there is a faint filling-in. The strength of the effect is similar to Test 4.

5.6.3 Discussion

The model recapitulates human data, but the strength of the effect is modest and much like Test 4.

5.7 Test 7: Static eyes over static line and background masked

The filling-in performance for Tests 3-4 could be considered modest at best. Here we present a solution that improves the results. In this test, we characterize how segregating the background into a different layer that is effectively masked improves the performance of filling-in the static line. It was found that perceptual completion must occur after similarity groupings and segregative processing have occurred (Kawabata, 1984) possibly related to figure/ground processing (R. J. Brown & Thurmond, 1993).

5.7.1 Methods

The methods used are identical to Test 4 except as detailed herein. The background segregation and masking technique are found on pages 98, 141, 211 and 274.

5.7.1.1 Apparatus

The software was engineered to inject the red horizontal bar stimulus into the appropriate photoreceptors. This was eye-position invariant.

5.7.1.2 Stimuli

With the equivalent stimulus of Test 4, but with the *person* not viewing any world, the code was instrumented to inject the horizontal red bar (3.6' height) onto the photoreceptors instead. Effectively, the dark background was masked so as not to interfere with the line in the fixational layer. This background segregation and masking

technique comes from pages 98, 141, 211 and 274. The injected LMS values equivalent to red=255 were (97, 50, 6).

5.7.1.3   Procedure

As in Test 4 (agent kept its eye static), except that the *person* was detached from the view and the stimulus was presented to the appropriate photoreceptors by instrumented code. This background segregation and masking technique comes from pages 98, 141, 211 and 274.

5.7.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s238). In this section, we summarize and discuss various highlights.

5.7.2.1   Blind spot filling-in behaviour

The model's behaviour of blind spot filling-in is summarized in Illustration 157.
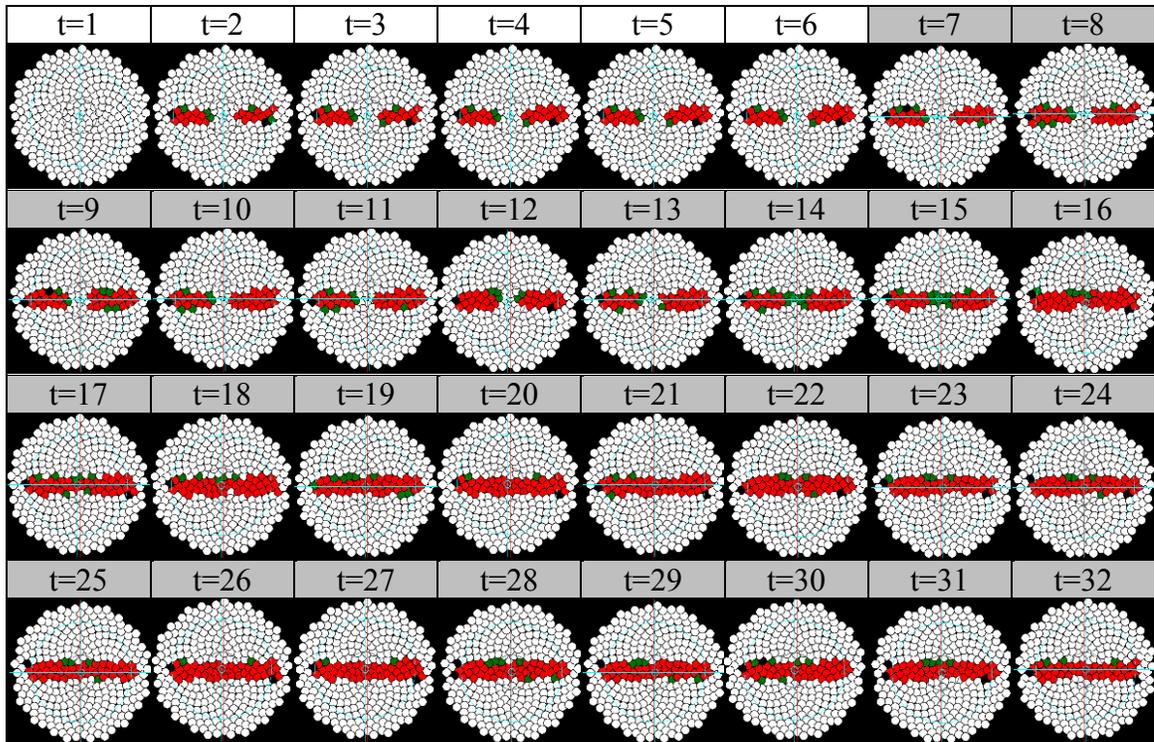
Illustration 157: Blind spot filling-in behaviour of 1ˢᵗ Level RFs with Static Line stimulus under mask

Similar spatiotemporal filling-in behaviour occurs at the 2ⁿᵈ and 3ʳᵈ RF levels. These

details are available in the supplemental material.

5.7.2.2   Filling-in performance details

Details of filling-in the 1ˢᵗ Level RFs is

shown in Illustration 158. Similar

spatiotemporal filling-in behaviour occurs

at the 2ⁿᵈ and 3ʳᵈ RF levels. These details
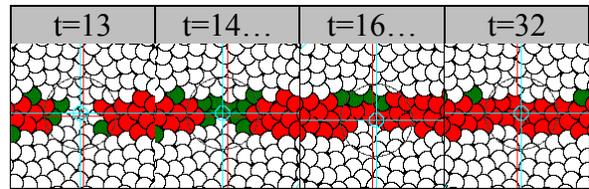
Illustration 158: Static blind spot filling-in details of 1ˢᵗ Level RFs

are available in the supplemental material. At tick 13, the blind spot was yet unfilled. At

tick 14, it filled-in with green much like Test 4. But the green disappeared and by tick 16,

we could unequivocally state that the horizontal bar was fully completed within the blind

spot. This gives a filling-in time of 14 ticks (from 2 thru 16). By tick 32, we demonstrate

that the filling-in process respects border contours and the redness does not diffuse into the rest of the blind spot. While the Test 4 results were weak, these results are strong.

5.7.3 Discussion

5.7.3.1 Why green filling-in again?

In Test 4, as in here, the horizontal line was first completed with green rather than red. We pinned the blame on photoreceptor and RF heterogeneity. However, we can clarify. Although we injected pure red into the eye, e.g., RGB values of (255, 0, 0), this was translated to the LMS values of (97, 50, 6). Thus, M only photoreceptor patches would appear as if they had an RGB value of (0, 119, 0) when converted backwards.

Therefore the green filling-in is a consequence of red filling-in with insufficient cone types. The green disappears once information about all cone types is combined within the RF. While an RF is a single Emergic Unit, each LMS value is kept as a separate Emergic Value pathway, and as such, their heterogeneities would differentially affect whether they are accepted by the Emergic Port of their destined unit.

5.7.3.2 Strength of filling-in

In Test 4 it was not clear whether filling-in completed as the RFs appeared at best just faintly reddish. Test 5 had slightly stronger result, but also suggested memory interference effects with the background. Indeed filling-in is ambiguous between the dark background and line foreground. Thus the RF statistics reflected the contribution of both dark background having large spatial extent, and the red foreground having small spatial extent. Therefore the amount of red was vastly reduced on the GUI. In this test, the

background stimulus was segregated to the background processing *layer* where it was effectively masked. Thus the RFs received only fixational layer information without interference. No matter how thin the red line would be, if fully detected by a single S/Red photoreceptor, then the higher level RFs would show the result in red at full strength on the GUI.

5.7.3.3   Time of filling-in

It took 14 ticks or nominally 140ms to fill-in the horizontal bar. However, blind spot filling-in is ordinarily thought of as instantaneous (Weil & Rees, 2011). Is this a problem in our model? Remember, this is 140ms for sensation, and who knows how much longer with conscious perception that we do not model.

Firstly, the qualitative "instantaneous" has been found to be greater than zero. The filling-in of a horizontal line has been found to take as long as 120ms (R. J. Brown & Thurmond, 1993; Kawabata, 1984). For monocular viewing where spatial filling-in must occur, there is a neural response latency of 44.8ms likely due to the slow horizontal signal velocities as compared to the fast feedforward or feedback velocities that can fill-in from the other eye (Matsumoto & Komatsu, 2005). Nevertheless peak neural responses still occur 80ms after stimulus onset. Neurally, we are in the correct ballpark.

Secondly, our tick rate has not yet been calibrated to biology. Thirdly, many aspects of the model have not been optimized such as the placement of RFs. In particular, an extremely large RF covering the entire blind spot would reduce our filling-in time substantially. However, it may have other negative consequences that we wish to avoid if

191

possible. The point of this Test 7 was to show the possibility of filling-in with large RFs. In particular, we wished to show that a layered segmentation process is an important architectural aspect – one should not try to optimize too early. Our intent is to induce the architectural pieces first, then optimize.

5.8   Test 8: Filled-in ring followed by saccade

While our model for filling-in the blind spot may be suboptimal and woefully incomplete in regards to the entire blind spot phenomenology, we wondered whether it could already clarify some, especially if they relate to filling-in under motion…

Shifting must be a mechanism behind the blind spot because it would otherwise be noticed during eye motion. It takes time for the blind spot to be filled-in under static eye conditions, but it is truly 0ms under dynamic eye conditions that include planned shifting. If the eye moved halfway through the static filling-in process, perhaps the partial filling-in could be perceived? In this test, we explain the results of (Spillmann, Otte, Hamburger, & Magnussen, 2006) as shown in Illustration 159 by our shifting mechanism.

Illustration 159: Partial filling-in of blind spot
From (Spillmann et al., 2006).

5.8.1   Methods

The methods used are identical to Test 1 except as detailed herein.

5.8.1.1 Subject analysis summary

This agent is similar to the one of Tests 4-6, but optimized to fit the ring stimulus. Moreover, its L:M:S cone ratio was made similar to humans as the test results will be compared against human data (Spillmann et al., 2006). This is summarized in Illustration 160 while details can be found in the supplement (s165). The visual

Illustration 160: Photoreceptor mosaic of agent used in blind spot tests 8

angle for the entire photoreceptor mosaic was 77.34′ while the visual angle for the 1st, 2nd and 3rd RF levels was 100.92′, 143.01′ and 113.34′ respectively.

5.8.1.2 Stimuli

In (Spillmann et al., 2006), the participant's blind spot was mapped, and then a thick contour formed the stimulus – they are often not circular. In this test, the blind spot was constructed as a perfect circle, and the ring stimulus was manufactured as shown in Illustration 161.
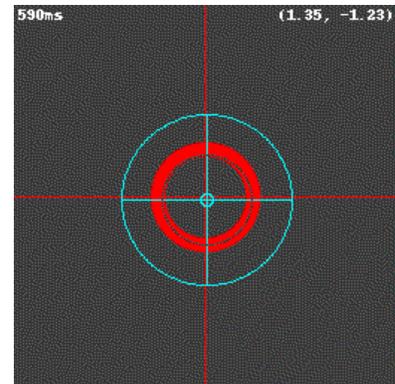
Illustration 161: Blind spot Test 8 stimulus (Ring) with overlay

5.8.1.3 Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to keep its eye steady (barring involuntary jitter) until the entire blind spot filled-in. At that time, the agent was instructed to saccade in the following direction and speed.

- Move right for 11 ticks at a speed of 3′/tick (± jitter), and at the same time
- Move up for 11 ticks at a speed of 2′/tick (± jitter)

The artificial brain was monitored for the last 32 time ticks. Note that there is an eye movement planning and distribution delay of 6 ticks before the eye actually moves as desired, so the above plan will not complete.

5.8.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s249). In this section, we summarize and discuss various highlights.

5.8.2.1 Blind spot filling-in then shift behaviour

The model's behaviour of blind spot filling-in for the 1st Level RFs is summarized in Illustration 162.
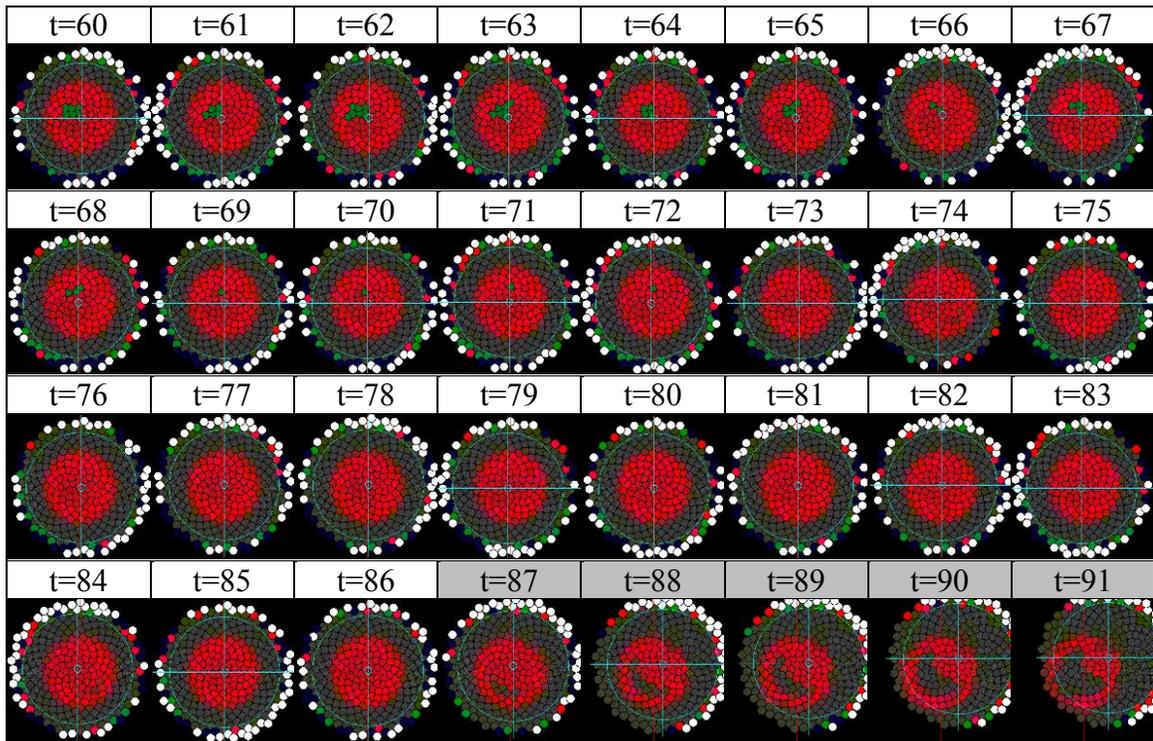
Illustration 162: Blind spot filling-in behaviour of 1st Level RFs with Ring stimulus

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

The blind spot was filled in by tick 73. They eye started to move at tick 87 which is indicated via a grey title background.

5.8.2.2   (Spillmann et al., 2006) redux details

We have recapitulated the partial filling-in phenomenon of (Spillmann et al., 2006), and this is shown for the 1st Level RFs in Illustration 163. Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.
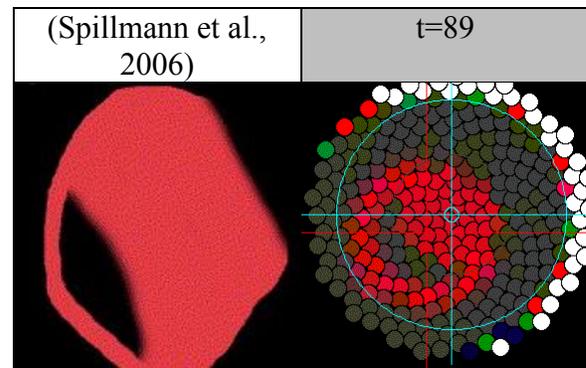
Illustration 163: Blind spot filling-in details of 1st Level RFs with Ring stimulus

5.8.3   Discussion

We have demonstrated that ECM can reproduce a novel blind spot phenomenon referred to as partial filling-in by (Spillmann et al., 2006). While Spillmann did note that this was due to eye movement, he supposed that eye movement prevented completion by "breaking overall adjacency". Our interpretation differs. The blind spot was completely filled-in, but the eye moved one way so that the background was shifted into the blind spot in the opposite direction. This is compatible with Spillmann's own observation that

> *"with improper fixation, one could actually see two dark "shadows", one on the side where the frame had invaded the blind spot and the other on the opposite side, where the frame had moved away from the blind spot border."*

While this observation begs for some shifting terminology, such is foreign. It is this shifting of information under eye movement that is ordinarily unaccounted by current

195

research. While it forms the basis for Tests 1-3, there is phenomenal precedent. For example, the motion of a moving object can be extrapolated into the blind spot (Maus & Nijhawan, 2008).

Spillmann exemplifies a monumental deduction problem in neuropsychological research in making sweeping overgeneralizations from conceptual models. He deduces:

> *"In general, the results are consistent with the recent model of (Komatsu, 2006), assuming a lateral spread of signals from oriented contrast sensitive cells along the edge of the cortical representation of the blind-spot region."*

Firstly, without a computational model, it is not clear that the term "consistent with" has been scientifically demonstrated beyond armchair handwaving that assumes simple linear combinations of isomorphic functions. What we show is that Spillmann's results have been scientifically demonstrated to be consistent with ECM. In particular ECM has no oriented contrast sensitive cells, no special cells along blind spot borders, and no active lateral spread of signals – we interpolate but do not extrapolate.

It is not that Spillmann's supposition is wrong. His results may be consistent with a plethora of multiple realizations. It is just that the supposition cannot help Spillmann in theorizing. It is up to the owners of the individual models (or realizations) to account for Spillmann's results, as we do. Moreover, even if Spillmann shows that his findings are inconsistent with a particular model; these models are constantly being improved. Models such as ours are non-linear with many emergic effects (border processing being an example), so what may not seem possible in a model today may be possible tomorrow

without overthrowing the basic premises of the model. Certainly model comparison requires a much more detailed analysis than provided by Spillmann.

In general, computational models are informed by data, but this is one example where the interpretation of data can be extremely informed by theory. We are not even suggesting that our theory is correct, just that there are multiple interpretations for "partial filling-in". Indeed, the descriptive name for the phenomenon "partial filling-in" is a poor choice as it already connotes a particular explanation and mechanism.

5.9 General results and discussion

It is widely believed that blind spot completion requires a mechanism that connects edges across space (Kellman, Yin, & Shipley, 1998). This would include the need to detect edges (or border contrasts). A few such conceptual models
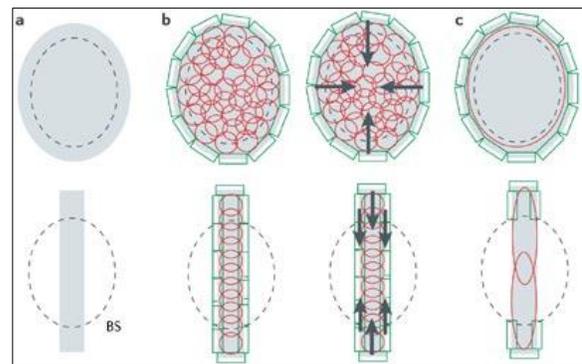
Illustration 164: Conceptual completion models
From (Komatsu, 2006).

are illustrated in Illustration 164 as reviewed by (Komatsu, 2006). ECM has no such mechanism, but border-like processing emerges due to the way *structured* information flows throw the system, i.e., because of the no extrapolation behaviour (p57). In particular, the information includes statistics over the spatial extent of surface colours. For example, if there was a horizontal red bar, then its statistical information could be for the R in RGB:

Value=255±0.0; x=0.0±radius of eye; y=0.0±1.8′

much like the stimulus of Test 7 on page 193. We should point out that the values are statistical distributions and the ±1.8′ represents one standard deviation, while the 0.0 represents the mean. It is the common statistics that connects one region to another over large RFs, and it is the "interpolate but do not extrapolate" behaviour (p57) that ensures information does not diffuse beyond emergent borders. The mechanism behind such emergic border completion is shown in Illustration 48 (p57). No matter how large an RF gets, it can still retain information about a thin line.

5.9.1 Huge RFs

ECM is a minimal unified model that aims to explain complex behaviour by re-using basic computational functions. We have only two such functions that interact at the locus of a single receptive field type. These are the *handle missing data* function and the *maintain information coherence* function. Without introducing either new functions or new component types, we hypothesized that the behaviour of filling-in the blind spot will be facilitated by having larger sized RFs. This hypothesis was supported in the model by Tests 4-8 and confirmed in Test 7. It finds support in biology as shown in Illustration 165.

ECM RF Levels (Tests 4-7)			Biology
1st Level RFs	2nd Level RF2	3rd Level RFs	(Komatsu et al., 2000)
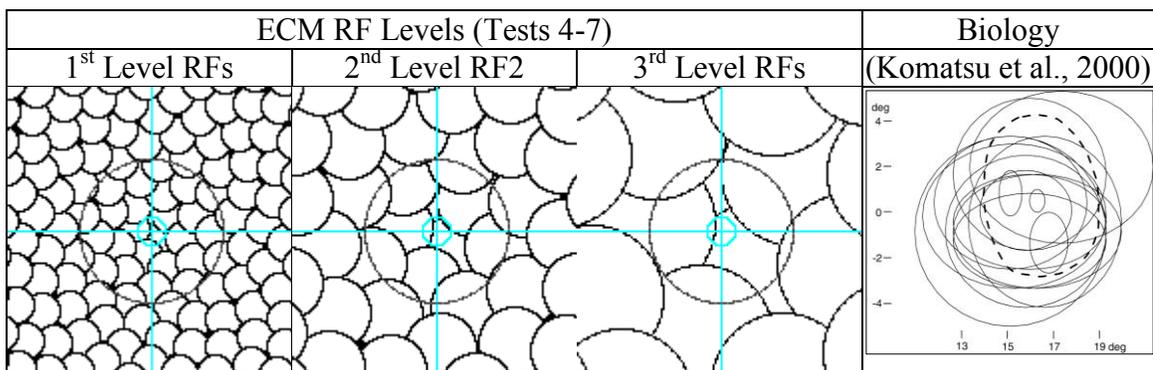
Illustration 165: Large RFs in ECM and biology

We should point out that the size of our RFs has not been optimized, and their placement is difficult to control due to the sunflower pattern development program.

198

There is a question about the function of these large RFs. If RFs are taken to be isomorphic to perception, then the large ones could only be long and thin to represent lines, as in column c) of Illustration 164. Moreover, such RFs would need to exist throughout the blind spot at all possible orientations. However, ECM provides an alternative explanation. RFs are not isomorphic but contain structured information that can flow and be interpolated thanks to dendritic processing. All possible orientations would not be required as that is contained within the informational content. Large RFs could allow for information integration and explain the better vernier hyperacuity over the blind spot than anywhere else (Crossland & Bex, 2009). Indeed, the large size of RFs within the blind spot supports our view that information is shifted within and across RFs.

A larger size simply makes it easier to move the information across large spatial regions, or maintain it internally. On reason to maintain internally may be to reduce dissipation, as shifting across similar sized RFs cause information to be split and merged because shifting cannot be in integral RF units or in the direction of repetition (there isn't any in the heterogeneity within ECM). Even in a perfectly homogeneous and uniform hexagonal mosaic, the direction of saccade, object movement or image contour would be restricted in alignment to at most 6 possible directions of symmetry in order to prevent dissipation.

5.10   Conclusion

We have demonstrated a model of filling-in the blind spot that is instantaneous under dynamic eyes and takes some time with static eyes (that nevertheless still jitter). Our

model is robust against a range of stimuli, and can reproduce novel phenomena such as partial filling-in, but with a computationally informed interpretation.

Our model has a minimal functional decomposition, and without introducing any new computations or neural types, we hypothesized that larger RFs over the blind spot could facilitate filling-in. This finds support both within the model and within neurobiology (Komatsu & Murakami, 1994). Large sized RFs only make sense if they manage information rather than be isomorphic to vision, and this is compatible with dendritic processing needed to handle information flows.

Appendices

Appendix A: An emergic model of filling-in for anorthoscopic perception

When a moving object is perceived behind a *narrow slit*, as in Illustration 166, it may appear whole despite being mostly obscured by the narrow *aperture* at any given time. This psychological phenomenon is called *anorthoscopic* perception and is likely the most

Illustration 166: Zoetrope replica
Image by (Dunn, 2004).

persuasive argument for having a detailed internal visual model of the world, despite "erroneous intuitions" (Simons & Ambinder, 2005) from change blindness studies (Rensink, 2002), or ecological/sensorimotor arguments for direct perception (Gibson, 1950; O'Regan & Noë, 2001).

In this section we present the world's first cognitive computational model of anorthoscopic sensation (supporting perception) based on the Emergic Cognitive Model (ECM) architecture (p323). While ECM does not have any specific mechanism dedicated to memory, nor to filling-in spatially distributed information across time, we hypothesized that both of its basic functions – *handling missing data*, and *maintaining information coherence* – could interact with experimental stimuli and measurements to cause this anorthoscopic filling-in phenomena to emerge as required. As part of characterizing the relative contributions of the two basic functions to the main phenomena, we also discovered that this model best demonstrates the emergence of

flowcentric information processing and highlights its superiority over the *neurocentric* accounts of today – accounts that implicitly model direct perception.

No other computational models have been proposed to explain this specific phenomenon, so we will contrast our solution to conceptual ones or to models of similar phenomena.

A.1 Introduction

Something like the zoetrope of Illustration 166 was invented in China around 180AD (Needham, 1962). The zoetrope exhibits several different kinds of visual phenomena. For example, like cartoon animations, several different image frames are fused together to bring about apparent motion. However, the specific phenomenon of interest to this chapter is where an image larger than can be seen through the narrow slit, is nevertheless perceived as a whole after it has fully moved past the slit. Many different scientists have historically been involved on variants of the zoetrope and their phenomenology (Wade & Heller, 1997). Joseph Plateau invented several, and in 1829 he first used the term *anorthoscopic* to mean abnormally viewed. Zöllner (1862) studied the phenomenon as *anorthoscopic perception*, but it has also been called *aperture-viewing* (Craddock, Martinovic, & Lawson, 2011) or mentioned in regards to viewing through a *narrow slit* (Öğmen & Herzog, 2010). It could be thought of as perception through an exceptionally large occlusion with holes (E. M. Palmer, Kellman, & Shipley, 2006).

Beyond perceiving the larger image, anorthoscopic perception includes artefacts such as image compression –
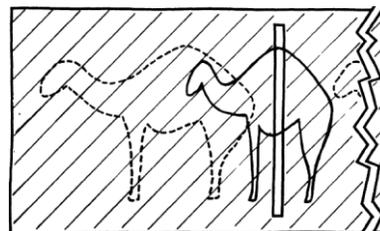
Illustration 167: Anorthoscopic camel passing through the eye of a needle.
From (T. E. Parks, 1965), Fig. 1.

a circle would appear as an oval – as well as spatial displacement (Sohmiya & Sohmiya, 1992). Parks (1965) is famous for distorting the shape of a camel as it passes "through" the eye of a needle (Illustration 167). We will focus our review on integrating the larger image from slices appearing within the narrow slit – ignoring the distortions and displacements for now.

A.1.1 Two forms of anorthoscopic perception and relation to motion

There are two forms of anorthoscopic perception (Mateeff, Popov, & Hohnsbein, 1993). As currently described, the narrow slit is fixed, and the image moves behind it. The opposite is also possible – the image is fixed and the slit is moved across the image. In this latter case, anorthoscopic perception is enhanced when viewers can actively control the aperture movement, possibly due to more accurate knowledge of speed, direction and extent of their own movement (Craddock et al., 2011). In this form, the phenomenon would be identical to wearing
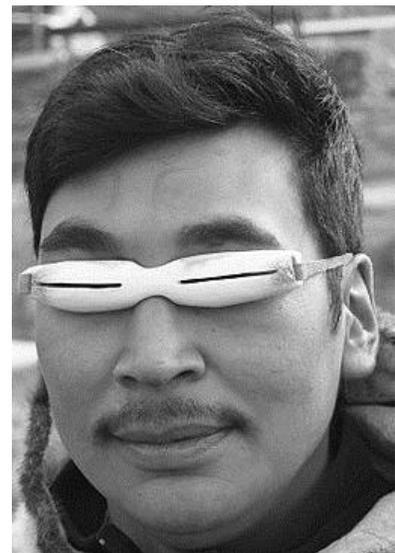
Illustration 168: Inuit snow goggles.
Image by (Idrobo, 2007).

snow goggles (as in Illustration 168) and perceiving the larger world by moving the head around. A 3D version of this experiment also demonstrates that active or *in-situ* manipulation of a disembodied eye (an ultrasound transducer) facilitates 3D perception while an *ex-situ* manipulation leads to more errors (Klatzky, Wu, & Stetten, 2010).

Similarly, if the eye is tracking a highly occluded and moving global object, then the eye motion facilitates the anorthoscopic perception of the object (Rock, Halper, DiVita, &

Wheeler, 1987). Unfortunately, this begs the question as to how the global object is tracked prior to its anorthoscopic perception. A good example is in Illustration 174 on page 212. It may be that only a few red portions of the fire truck provide unambiguous motion signals, but the entire fire truck could be loosely bound by its common redness. Once tracking starts, then the eye motion could facilitate the anorthoscopic perception. This is simply our supposition on unaddressed matters. Motion processing is known to be complicated and involving both bottom-up and top-down process (Whitney, 2006), and these must occur before or jointly with anorthoscopic integration. Kandil & Lappe (2007) show that when local motion information is not available, it may be inferred from global information across several occlusions.

The eyes maintain a dynamic representation of object motion so they can track it even when hidden to both (Orban de Xivry et al., 2008), and this allows accurate depth perception even if the object is always invisible to one of the eyes as in Illustration 169.
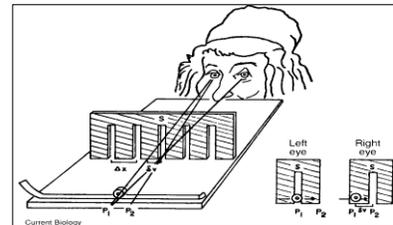
Illustration 169: Depth through picket fence.
From (Burr & Ross, 2004).

The existence of these two presentation forms has not been generally appreciated. When the eyes are fixed, it implies two distinct cognitive processes. Thus, many of the empirical arguments may be faulty in assuming just one phenomenon and one mechanism. Indeed, mechanism and presentation may be conflated. Morgan, Findlay, & Watt (1982) argued for two mechanisms depending on whether the eye moves slowly, or if the eye moves quickly to allow for retinal painting.

Of course, both forms can occur simultaneously and this can lead to Roget's illusion (Hunt, 2003). When a rotating wheel is seen through a moving set of slits or Palisade, the spokes will appear curved as shown in Illustration 170.
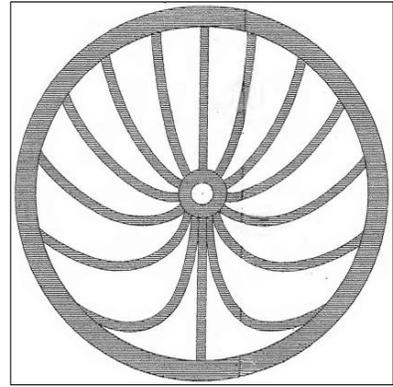
Illustration 170: Roget's Palisade From (Roget, 1825).

A.1.2   Retinal Painting

There has been one hypothesis that did not involve completing an internal model of the world from image slices. In 1867, Helmholtz suggested that the changing view within the narrow slit was "painted" across the retina as the eyes moved in parallel with the image behind the slit (Helmholtz, 1925). This became known as the inpainting hypothesis. This hypothesis happens to be a form of direct perception (no internal model or representation of the world is required), and thus, *inpainting* is a common explanation to resolve other aspects of direct perception such as those surrounding optic flow fields (Berkels, Kondermann, Garbe, & Rumpf, 2009).

There are many arguments against inpainting and for post-retinal processing. Firstly, Parks (1965) noted that if two figures pass simultaneously under the aperture but in opposite directions, then they will both be perceived as complete (and having opposite distortions). However, a single eye cannot move in two directions at the same time to be appropriately painted.

Secondly, the inpainting hypothesis depends on the persistence of vision to *hold the paint*. However, this was debunked by Wertheimer (1912) who linked the illusion of persistence to motion processing. Indeed, anorthoscopic perception cannot occur without motion. Nevertheless, the myth of persistence continues today not only in film studies (J. Anderson & Anderson, 1993), but implicitly throughout the visual perception literature whenever any kind of visual short term memory or icon (e.g., McCloskey & Watkins, 1978) is mentioned.

Thirdly, Dixon & di Lollo, (1994) noted that any form of intrinsic persistence would lead to image smear of moving objects (see also Öğmen & Herzog, 2010). We will review image smear problems in greater detail in the next chapter. Dixon & di Lollo proposed a stimulus coding strategy that would allow for both integration and segregation demands to be met in competing phenomena – the segregation would prevent integrative smear and provide motion change information. However, their proposal was not computational, not demonstrative, and therefore not persuasive. Similarly, Palmer et al. (2006) hypothesized *dynamic visual icons* that we suspect could be immune to smear, but again, no computational model. The numbers of conceptual toothbrushes just multiply and with minimal support, one wonders whether they should be granted the status of *theory* (Gigerenzer, 2010). It is also no wonder that the allure of persistence remains – it is a simple explanation covering simply all phenomena and therefore simply too powerful to resist – even if entirely misguided. Compare to:

> *"We suggest that these [non-retinotopic representation] manifolds are created by the emergence of dynamic reference-frames that result from motion segmentation.*

We also suggest that the metric of these manifolds is based on relative motion vectors.” (Öğmen & Herzog, 2010)

Even if correct, it is difficult to comprehend and easily dismissed. Moreover, *suggestions* seem to be synonymous with toothbrushes. The myth of persistence has been living now past its due date for 100 years!

Fourthly, anorthoscopic perception can occur in the absence of eye movements or when the image is stabilized on one part of the retina to prevent inpainting (Fendrich et al., 2005; Morgan et al., 1982).

A.1.3   Time of arrival hypothesis

An alternative to direct perception is the postretinal time-of-arrival hypothesis positing that the stimulus is reconstructed to its extended spatial form based on the time of arrival of each slice (T. E. Parks, 1965). Such a theory is insufficient (Öğmen & Herzog, 2010) as
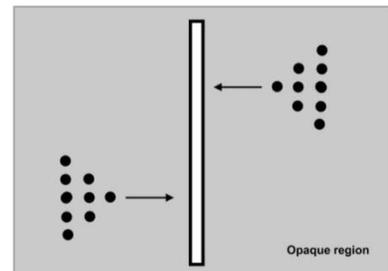
Illustration 171: Underdetermined time-of-arrival stimulus
From (Öğmen & Herzog, 2010).

demonstrated in Illustration 171. In this case, two different images cross the narrow slit providing identical time-of-arrival information, yet their direction of movement, and orientation differ. Thus, the motion of the image within the aperture must be taken into account (Sohmiya & Sohmiya, 1994). Motion information has been found to be critical to integrating and obtaining clear percepts (Nishida, 2004). In the example of Illustration 171, if the slit was exactly one pixel wide, and the objects were moved one pixel at a time, then even local motion signals would not be available. However, with multiple slits, the global cues can provide for the missing motion direction (Kandil & Lappe, 2007).

A.1.4 Position updating hypothesis

Palmer et al. (2006) suggested that the velocity of hidden objects would be retained and used to update their hidden position. It requires segregating the image layer into a persistent medium that allows the piecing together of object parts as shown in Illustration 172. This conceptual model is disembodied effectively with stationary eyes looking straight ahead.
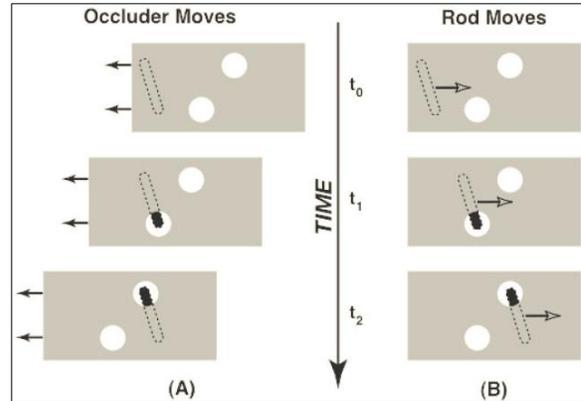
Illustration 172: Position updating hypothesis {COPYRIGHT}.
From Palmer et al. (2006), Fig. 3. Rod is retained in persistent layer. A: occlusion only moves and persistent medium updated with whatever appears. B: rod moves and its position is updated.

In reality there would be an additional four more cases, much likelier than the first two. These would be with either the rod or occluder moving, combined with either the eye tracking the rod or the eye following the occluder.

This theory is similar to the piecing together of sparse representations within a *schematic map* (Hochberg, 2007).

A.1.5 Current state of the art

The conceptual state of the art can be summed up as follows:

> *"Despite more than 140 years of research, the role of eye motions, retinal painting, and retinotopic storage in the formation of anorthoscopic figure percepts and their role in determining the characteristics of these percepts have remained unclear." – Rieger, Grüschow, Heinze, & Fendrich (2007)*

Therefore, it is not surprising that the causal and theoretical state of the art can be summed up as:

> *"The mechanisms underlying anorthoscopic perception are poorly understood." – Öğmen (2007)*

We could not find any computational model for anorthoscopic perception.

While all related research has involved hard experimental work, *theory* in its proper sense (Gigerenzer, 2010) has been scant. Perhaps it is time to stop armchair conceptual theorizing and start computational modeling. We hope to show the value in doing so. In particular, we hope to show that the conceptual work within experimental research must walk hand-in-hand with theoretical development. Theory informs the concepts behind empirical research. Without theoretical development, empiricism can become stuck with underpowered and incommensurable conceptions. In particular, we want to introduce the concept of *flowcentric* information processing.

A.2   Emergic Model

In order to explain the phenomena of this chapter, we have taken the Emergic Cognitive Model (ECM) architecture as summarized on page 5, and made no changes. For anorthoscopic perception, we have two intuitions as to how ECM might explain the results.

A.2.1   Identical to trans-saccadic and peri-saccadic perception

Firstly, the system is structurally identical to "An emergic model of filling-out for trans-saccadic integration" (p60). However, rather than filling-out the extra-retinal region by shifting information from within the retinal region, we fill-in the retinal region that is masked beyond the narrow slit by shifting information that is not masked within the slit. *Filling-in*, *filling-out* or *shifting* become synonymous terms. The important aspect is that shifting occurs by updating the internal model of the world to compensate for eye motion.

In the future, such compensation will also take into account object motion, but for now, only self-motion is accounted for by the *maintain information coherence* function.

A.2.2 Black vs. suppression

There is one significant difference between this model and that of "3 Chapter:" (p60). In the latter, the extra-retinal region does not receive any visual information bottom-up from the photoreceptors as it is beyond their reach. The *handle missing data* function will therefore make use of the lateral shifted information. However, in this chapter, the region that is masked beyond the slit should technically still receive a bottom-up colour value of black. If this actually occurred, then lateral information would not be used, there would be no integration of information, and the phenomenon of anorthoscopic perception would not occur in our model. How can we rationalize the suppression of black? The answer we sketch here is part of our theory of visual processing, and even though it is not yet implemented, it does allow for our apparent slight-of-hand.

A.2.3   Segmentation layers

ECM has a larger theory of visual processing as schematized in Illustration 173. It augments the implemented account with at least three tentative image segmentation layers. Occlusions and transparency information would be routed to the *overlay* layer, while information beyond fixation would be routed to
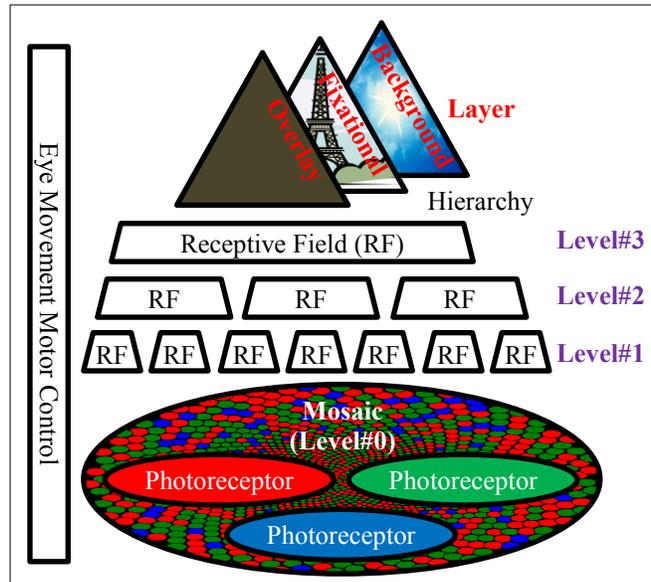
Illustration 173: Schematic of ECM with layered segmentation

the *background* layer. Visual information can be dynamically re-segmented with a change of attention. None of this layering mechanism is implemented, so the current ECM architecture can be thought of as instantiating the single *fixational* layer of visual processing (*layers* and *levels* are orthogonal). However, the black information that results from the narrow slit would be routed to either the overlay or background layer leaving no bottom-up information flows to the fixational layer. Therefore, our account of treating visually masked information as suppressing bottom-up flows of information into the fixational layer of visual processing is in line with our theory. However, the terms information *suppression*, *masking* or *omission* are not valid in our account as all visual information, such as black, is never thrown away, but simply dynamically routed where required. After all, it is possible to attend to the black mask or blink (O'Regan & Noë, 2001). With a computationally informed information processing account, such terms

become descriptive relative to a restricted point of view – the attentional layer in this case.

This same rationalization will be used in the subsequent phenomena of brief stimuli presentations, and eye-blinks. It is also plays an unexplained part of the eye paralysis phenomena of the previous chapter.

Our layered theory is not unique. Palmer, Kellman, & Shipley (2006) present a completely non-computational theory whereby the layer to be completed is determined by common object motion and this allows for the stitching together of jigsaw pieces over time as shown in Illustration 174 (and Illustration 172). This example has the benefit of providing an ecologically relevant situation, although we prefer a monkey swinging through the trees example to provide for an evolutionary plausible story – perhaps too, a lion hiding behind tall blades of grass with only a few narrow slits revealing portions. In any case, our account has significant aspects determined by computational code.

Illustration 174: Accumulating moving fragments.
From Palmer, Kellman, & Shipley (2006).

A.3 Test 1-3: Robust stimuli

In these tests, we demonstrate the robustness of our anorthoscopic filling-in model by examining its emergic behaviour across three different stimuli.

A.3.1 Methods

The methods used are identical to those of page 80 except as indicated herein.

A.3.1.1 Subject analysis summary

The details for the virtual agent can be found in the supplementary material (s260). However, the software was further instrumented to neutralize all the photoreceptors except for a central 8′ wide strip as shown in Illustration 175. The visual angle for the entire photoreceptor mosaic was 88.51′ while the visual angle for the 1st, 2nd and 3rd RF levels was 98.86′, 104.18′ and

Illustration 175: Photoreceptor mosaic for tests 1-4.
Suppressed photoreceptors are shown in grey to mimic anorthoscopic perception.

110.59′ respectively. Unlike the photoreceptor mosaic, the RF levels did not have any portions neutralized as they were used to demonstrate anorthoscopic perception.

A.3.1.2 Stimuli

Similar stimuli as in tests 1-3 of Leibovitz (p60) were used and are shown in Illustration 176, Illustration 177 and Illustration 178. They are fully described in the supplementary material (s262, s263).
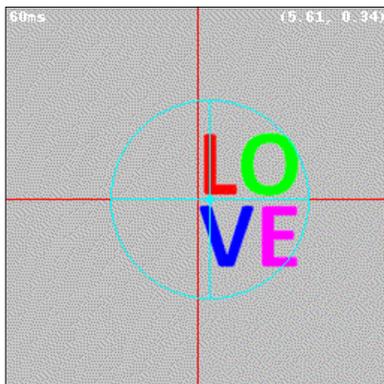
Illustration 176: Test 1 stimulus (Love) with view overlay
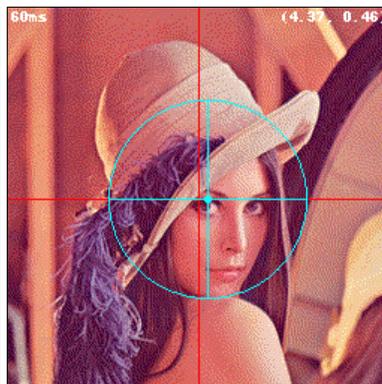
Illustration 177: Test 2 stimulus (Lena) with view overlay
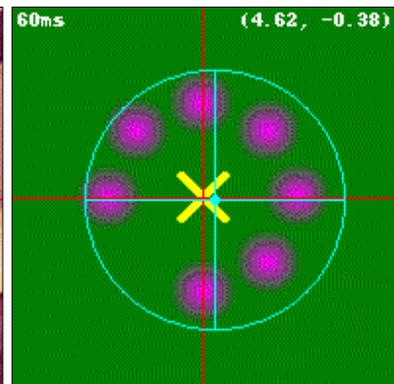
Illustration 178: Test 3 stimulus (Chaser) with view overlay

A.3.1.3 Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to move its eye in the following *Cylon* pattern (Bilder, 2009).

- Move right for 8 ticks at a speed of 5′/tick (± jitter)

- Move left for 8 ticks at a speed of 5′/tick (± jitter)

- Move right for 8 ticks at a speed of 5′/tick (± jitter)

- Move left for 8 ticks at a speed of 5′/tick (± jitter)

The artificial brain was monitored from tick 6 thru 37 as it takes 6 ticks before motor plans take effect.

A.3.2 Results

The visual frame-by-frame response results with precise timing information are detailed in the supplementary material (s265, s276, s287). When presented as an animation (Leibovitz, 2012a), all the results for this chapter can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=434
```

In this section, we summarize and discuss various highlights. We do not extract quantitative results as we merely wish to demonstrate anorthoscopic behavior. We also do not characterize behaviour across a range of speeds and cognitive fan-out factors as this would be premature considering the current simplicity of the ECM architecture – it is too early to engage in optimization activity.

A.3.2.1 Anorthoscopic filling-in behaviour for Level 1 RFs

While the narrow slit is 8′, the eye has moved back and forth twice across a 40′ field of view, and built up and maintained an anorthoscopic memory via spatiotemporal filling-in. This behaviour is demonstrated in Illustration 179, Illustration 180 and Illustration 181.
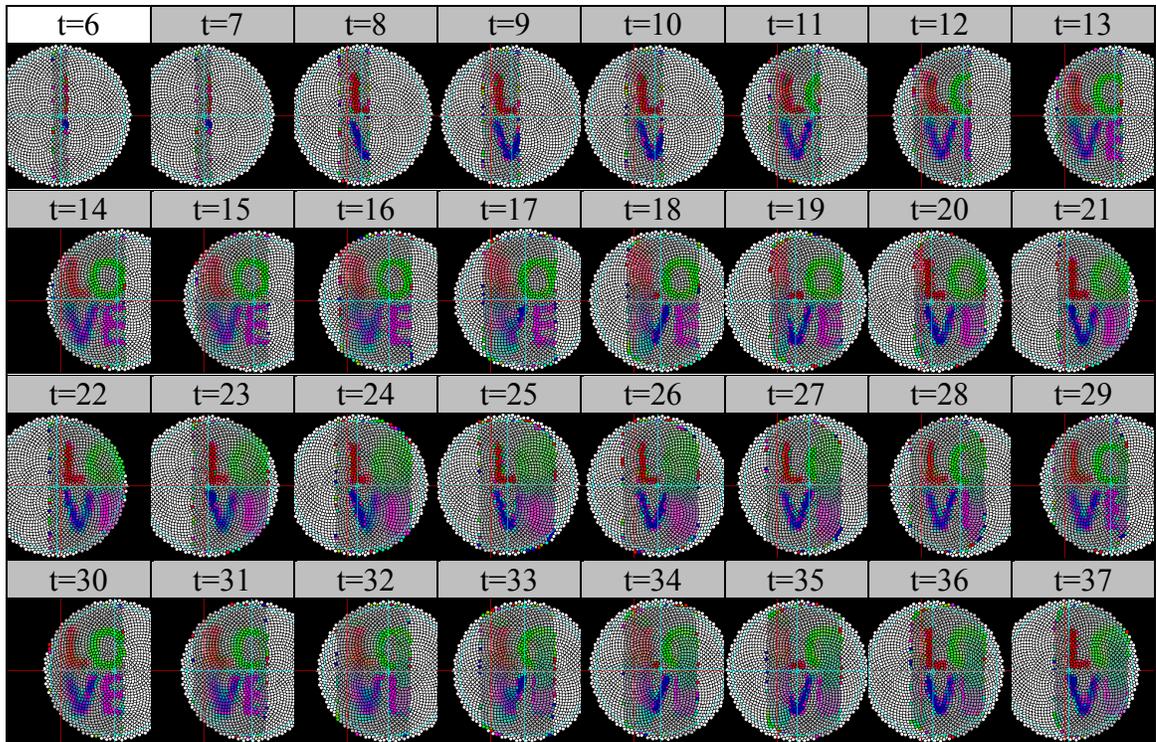
Illustration 179: Test 1 (Love) anorthoscopic filling-in behaviour
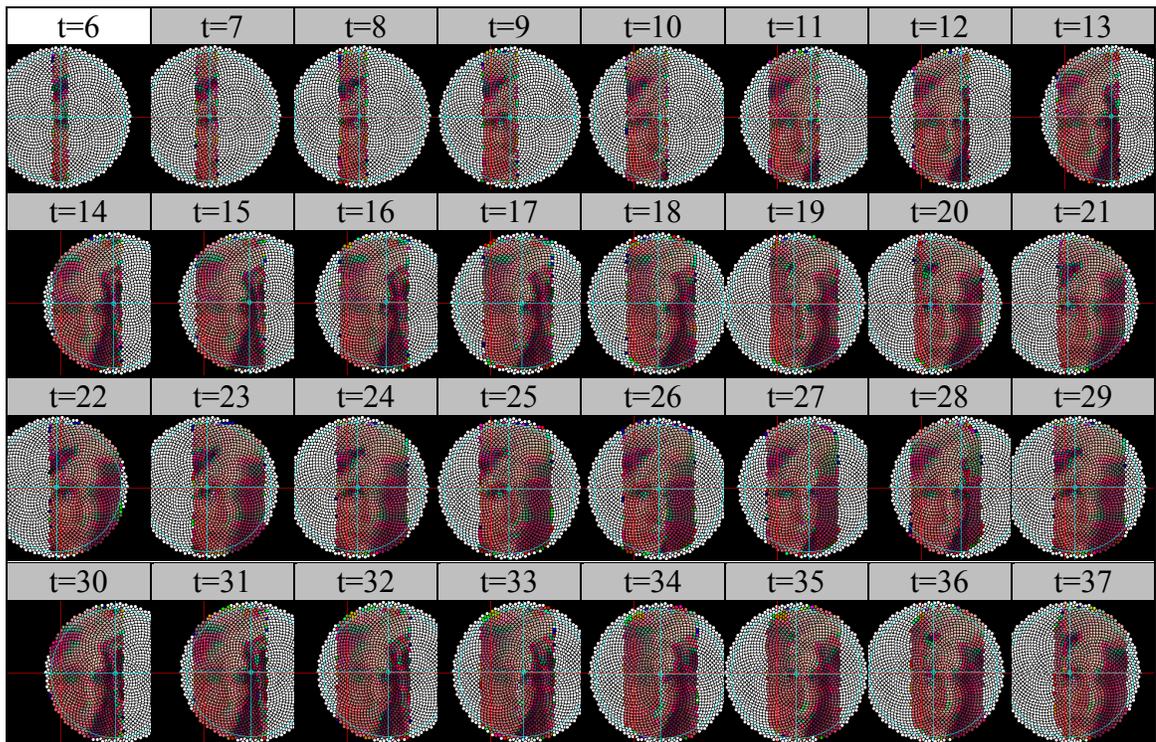
Illustration 180: Test 2 (Lena) anorthoscopic filling-in behaviour
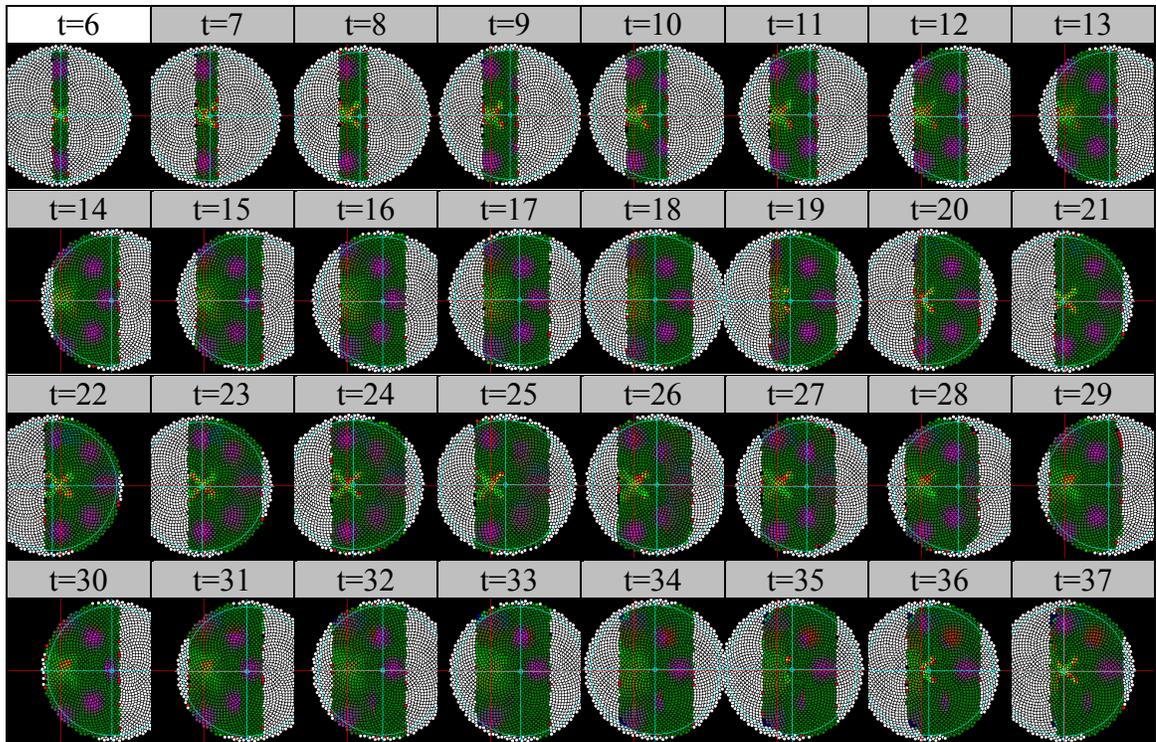
Illustration 181: Test 3 (Chaser) anorthoscopic filling-in behaviour

Similar behaviour occurs at the 2nd and 3rd RF levels.

Similar behaviour occurs at the 2nd and 3rd RF levels.

216

A.3.2.2 Anorthoscopic filling-in details for Level 1 RFs

For test 1, the best anorthoscopic memory example occurs at tick 30, while the worst at tick 34 (Illustration 182). Details for test 2 (Illustration 183) and test 3 (Illustration 184) are similar as is the case for the 2^{nd} and 3^{rd} level RFs.

t=30 (Best)	t=34 (Worst)
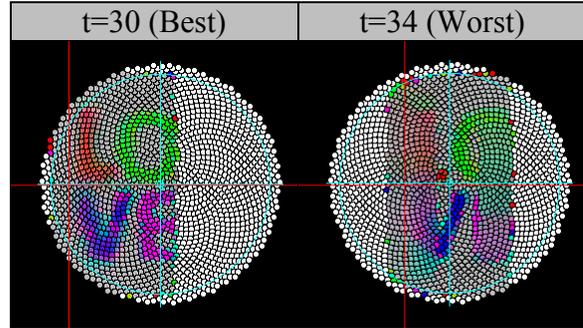
Illustration 182: Test 1 (Love) memory details

A.3.3 Discussion

A.3.3.1 Robustness

Many models work well for simplified stimuli but not natural ones (Kirlik, 2007). Here we demonstrate that the ECM architecture is robust against stimuli kind and without any change of parameters – it has no free parameters. Indeed, performance on natural images is better than letters in agreement with "3 Chapter:" (p87).

t=21 (Best)	t=34 (Worst)

Illustration 183: Test 2 (Lena) memory details

t=23 (Best)	t=34 (Worst)
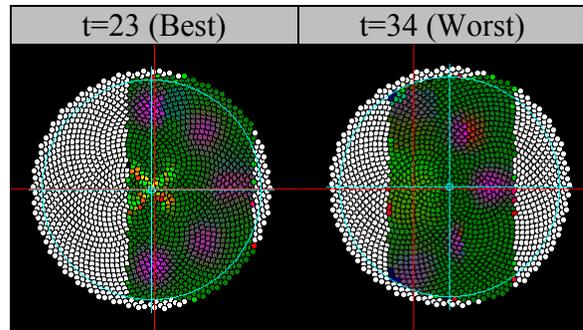
Illustration 184: Test 3 (Chaser) memory details

A.3.3.2 Performance

However, our model is far from perfect. In the worst case (test 1), the letter 'L' has lost cohesion due to interference with the background. At tick 34, it had gone for 8 ticks without refresh. This dissipation is already known (p86). In the future, when the grey background is segmented into the *background layer*, we expect that the memory

performance of the *fixational layer* to improve. Further chapters in this thesis will demonstrate that memory performance can be very good in the right conditions. This was already hinted at in the "drawing figures" test under eye paralysis (p98).

A.3.3.3 Flowcentric memory

The Lena stimulus is ideal to demonstrate the *flowcentric* nature of our information processing approach. In Illustration 185 one can clearly see that the same 40′ wide

| t=14 | t=22 |
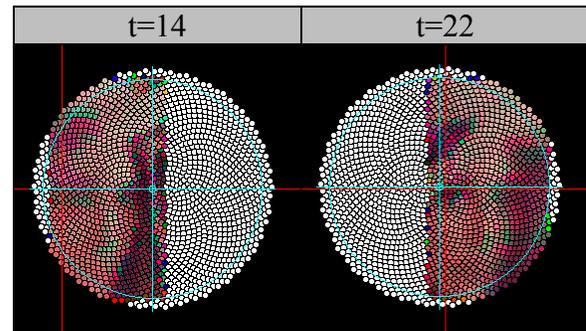
Illustration 185: Flowcentric memory

memory of Lena at tick 14 appears on the left of the retina, while at tick 22 it appears on the right. This will also occur at the 2nd and 3rd RF levels. The point is that the full memory of Lena is not associated with any neurons either in the retinotopic or spatiotopic reference frames. Certainly the 8′ *direct* sensory slice is retinocentric, but it is highly transitory and does not constitute an *indirect* memory using the terminology of direct perception (Gibson, 1950).

The question is where does the indirect memory of Lena reside, and is *memory* even a valid construct? To answer the first question, the memory of Lena resides in the continuous flow of information. Lena is perceptually stable within the flow as indicated by the red axes. One can make an analogy to the flow of water in a river. One can drop some coloured oils onto the surface of the water, painting a simple picture if you will, and this picture will maintain cohesion for some time. It does not really matter if the river turns left or right, or if it even narrows and moves more quickly. Within the flow of water, the picture maintains its relative characteristics. Of course, the picture will

218

continuously change in an unexpected manner, and in this sense, perhaps it should not be called a memory. This would be unlike *neurocentric* memories where the image might become fainter, but will not lose its spatial precision.

The purpose of ECM's *maintain information coherence* function is to ensure that information processing within the brain/mind maintains flow coherence throughout regardless of the current set of neurons keeping the flow moving. We think of neurons as paddle wheels keeping the water *computationally* flowing – it does not really matter where these water wheels are placed. This analogy is similar to the internet. It is the relative flow of packets that is important regardless of which routers they pass through in order to get to their destination. It does not matter too much how many routers there are (beyond delay times), where these routers are, and whether each packet follows the same route. This does not mean that one shouldn't analyse routers, but that an analysis of the data flow may be equally important.

Conceptually, at the boundary of sensors and effectors, the internal flow of information is used to engage with a flow of external interactions. In this sense the flow is always spatiotopic.

A.4   Test 4: Simple object (Square)

We demonstrate the robustness of our model by examining its behaviour across a wide range of stimuli. For this test, a simple object (square) will be used. This stimulus was chosen as it will be compared to its black and white counterpart demonstrating smear (Grossberg & Mingolla, 1993) here, and smear reduction (Francis & Grossberg, 1996).

The smear reduction comparison must await Leibovitz (p103). This stimulus does highlight an underdetermination problem of perception across small apertures.

A.4.1 Methods

The methods used are identical to Test 1 except as detailed herein.

A.4.1.1 Stimuli

A simple square object stimulus was used as shown in Illustration 186. It is fully described in the supplementary material (s263).
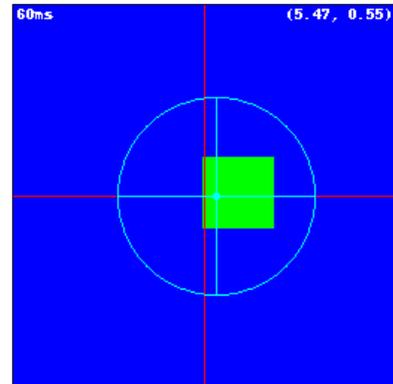
Illustration 186: Test 4 stimulus (Square) with view overlay

A.4.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s298). In this section, we summarize and discuss various highlights.

A.4.2.1 Anorthoscopic filling-in behaviour for Level 1 RFs

While the narrow slit is 8′, the eye has moved back and forth twice across a 40′ field of view, and built up and maintained an anorthoscopic view via spatiotemporal filling-in. This behaviour is demonstrated in Illustration 187.

t=6	t=7	t=8	t=9	t=10	t=11	t=12	t=13
t=14	t=15	t=16	t=17	t=18	t=19	t=20	t=21
t=22	t=23	t=24	t=25	t=26	t=27	t=28	t=29
t=30	t=31	t=32	t=33	t=34	t=35	t=36	t=37
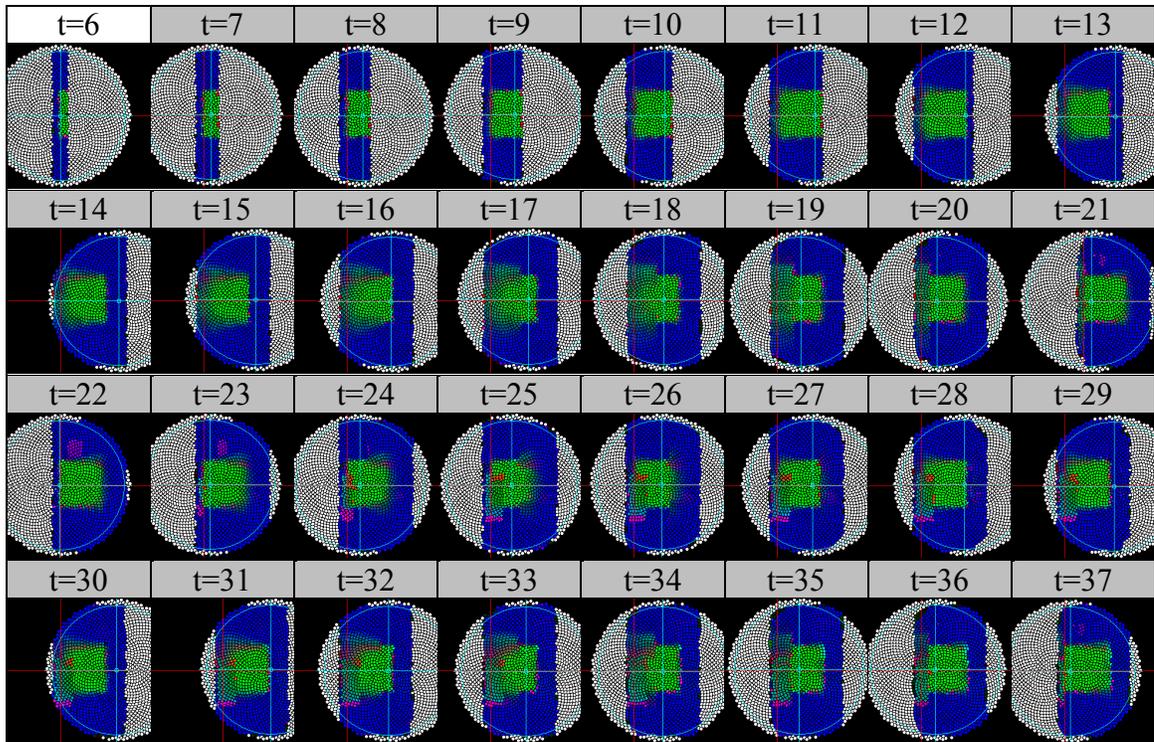
Illustration 187: Anorthoscopic filling-in of Level 1 RFs with simple object (Square)

Similar behaviour occurs at the 2nd and 3rd RF levels.

A.4.2.2   Anorthoscopic filling-in details for Level 1 RFs

The best anorthoscopic memory example occurs at tick 22, while the worst at tick 34. These are shown in Illustration 188. Similar examples can be found for the 2nd and 3rd level RFs.

| t=22 (Best) | t=34 (Worst) |
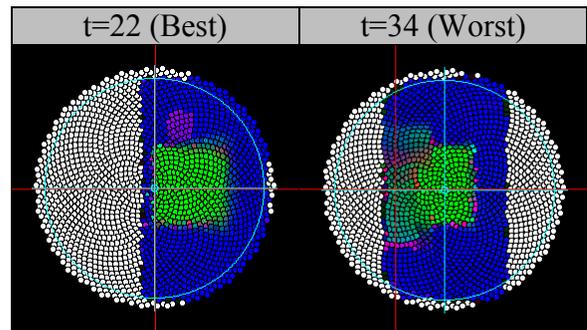
Illustration 188: Test 4 (Square) memory details

A.4.3   Discussion

The ECM architecture does demonstrate anorthoscopic perception to the square stimulus. The performance is similar to that of letters spelling LOVE. We are not looking to prematurely understand nor optimize such behaviour.

A.4.3.1 Underdetermination

Because the stimulus is spatially regular and has no texture, it becomes difficult to determine motion across a small aperture.
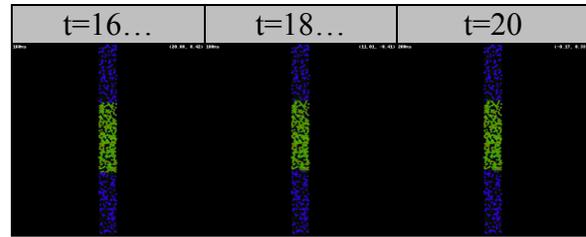
Illustration 189: Underdetermination of motion

For example, the photoreceptor inputs between ticks 16 thru 20 are identical as shown in Illustration 189. This is perhaps one reason why anorthoscopic perception is better with free viewing (where eye motion is known) then with fixed eyes with object motion that can be underdetermined (Rock et al., 1987). It is also why conceptual theories such as Palmer, Kellman, & Shipley (2006) mentioned on page 212 cannot be taken seriously if it has no computational backup – many of the fire-truck snippets through the tree branches would be underdetermined and could not be used to segregate objects into a layer based on common motion. Non-computational "theories" are a dime a dozen. Minsky famously thought that all of machine vision could be developed by a student as a summer project (Cosmides & Tooby, 1994).

A.4.3.2 Image smear: ECM vs. FACADE

In this section, we compare the effect of motion on image smear within ECM and FACADE (Grossberg & Mingolla, 1993).

A.4.3.2.1 ECM

The photoreceptors move across the right edge of the square during ticks 10 thru 13 as shown in Illustration 190.

| t=10 | t=11 | t=12 | t=13 |
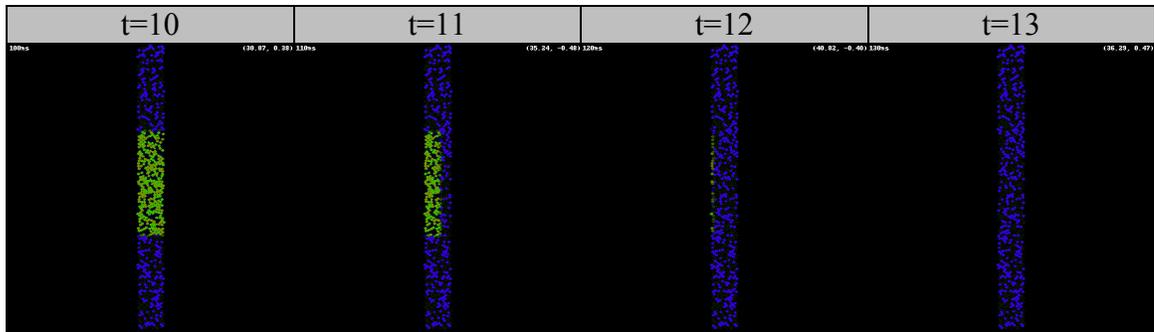
Illustration 190: Photoreceptors across edge

However, despite such motion, the corresponding edges within the 1st level RFs do not appear smeared during ticks 11 thru 14 as shown in Illustration 191.

| t=11 | t=12 | t=13 | t=14 |
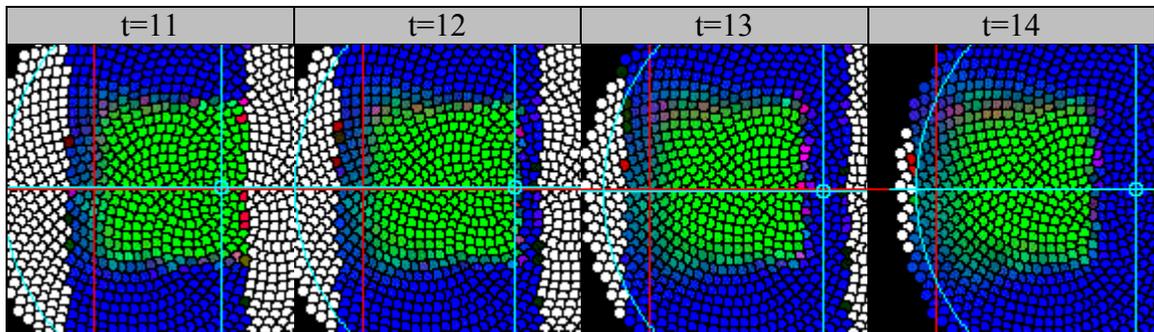
Illustration 191: No smear within RF layers

Similarly, the other RF levels also do not exhibit image smear.

A.5 General results and discussion

We have demonstrated a robust form of anorthoscopic perception whereby a stationary stimulus is effectively scanned by an eye and a narrow aperture that jointly move across. In such a form, the retina could not be painted and an indirect model of the world is required. The locus of the internal model is within the flow of information and this gives it properties unlike traditional conceptions of memory. The other form of anorthoscopic perception occurs when the aperture is stationary and the stimuli move underneath. That form is more complex and can be underdetermined, so we leave it for future research.

A.5.1 Flowcentric information processing

In a causal *flowcentric* account of short-term information processing, neural groups perform an interactive computation and transmit their values. They do not retain these values as memories, and it is not even a valid question as to whether such neurons are retinotopic or spatiotopic. Information can be locally retained within a flow of information, if there happens to be lateral connectivity among the neurons that keep the flow moving. One of the functions of neural computation is to maintain the coherence or referential integrity of causal flows despite eye movements. This ensures that the internal causal flow can be considered as spatiotopic and one with the external causal flow. However, an ephemeral flow would be a challenge to measure with traditional neurobiological techniques as it requires relative dynamic correlations on physical neurons – it looks so much better visually. Such techniques are just being developed (e.g., Ganmor, Segev, & Schneidman, 2011; Sporns, 2011; van den Heuvel & Hulshoff Pol, 2010) but a computational account will still be required to bridge brain with behaviour (Carandini, 2012).

While long-term storage is certainly distributed within neural and synaptic structures, the moment these memories are accessed they become part of the short-term information processing flow.

A.5.2 Robust and parameter free

The computational model is robust and parameter-free against a range of stimuli. Indeed, it is very good with natural images. It does not need a parameter to indicate the magnitude of motion.

A.5.3 Information suppression, omission or masking?

The narrow slit can be described as transmitting colour image information within the slit to the retina, and black elsewhere to the retina. ECM has an unimplemented theory of visual processing that would route and segregate the black information to an occluding/overlay visual processing *layer* for further processing. One could, after all, decide to attend to the black occlusion, but for our purposes, it is cognitively ignored. The fixational *layer* does not receive any colour information. In ECM, not receiving allows the use of lateral information flows or "memory" and it is this behaviour that leads to the integration of the image over space as the eye moves.

From the perspective of the entire cognitive system, the black is processed and one could decide to attend to it. From the perspective of the fixational layer alone, it would appear that the black is *masked*, *supressed* or *omitted*. A full computational information processing account simply highlights the fact that *mask*, *suppression* and *omission* are not useful causal terms for explaining the behaviour of an entire cognitive system. We use them interchangeably for descriptive effect only.

A.5.4 Eye movement

ECM explains the paradoxical finding that eye motion improves anorthoscopic perception even though there is neither net motion, nor painting across the retina (as with wearing snow goggles). Simply that self-motion is precisely known and effects perfect shifting. In the other form of anorthoscopic perception, object motion is often underdetermined and requires additional global integration if available.

While our anorthoscopic form is over a fixed image, Brockmole & Irwin (2005) have made an interesting hypothesis for the dynamic case of general integration.

> *"In the dynamic case, however, VSTM was generated in a location that did not overlap with one of the to-be-integrated percepts. In order to "move" the memory trace, overt shifts of attention were required."*

One could think of our model as a computational validation of their intuitions.

A.5.5   Position Updating

The position updating hypothesis (E. M. Palmer et al., 2006) is closest to ours in three respects. We both update the position of hidden object; we both treat the problem generically (arthroscopic perception is nothing more than a large occlusions with a hole); and we both require a layered segmentation to separate occlusions from objects.

We also differ in three respects. While their model is conceptual, ours is computational and demonstrative. While their model is disembodied, geometrical, and not cognitive, ours is cognitively embodied and takes eye motion into account. Additionally, ours is *flowcentric* rather than neurally spatiotopic.

Thus we differ in what gets updated and when. Palmer does not update the disembodied position of anything when the occluding slit moves (they say nothing about eyes). They do not update anything if the object remains stationary. We, on the other hand, shift cognitive information from neuron to neuron as the eyes move over a stationary object. Despite this, the non-physical flow remains spatiotopic. Thus we have opposite predictions. At least we make and demonstrate a neurocognitive prediction that can be empirically tested.

A.6 Conclusion

A.6.1 First to model and demonstrate support for arthroscopic perception

We have demonstrated the first computational model of arthroscopic *sensation* where joint eye and slit movement assists in filling-in the hidden details of a large object that can only be sensed piecemeal. Our models specifically support, but do not account for, conscious *perception*. Our "memories" only need to last long enough as to be perceived.

A.6.2 First to model and demonstrate a detailed internal visual model of the world

Our model has unique aspects and properties. It is the first computational account that *maintains* a detailed internal visual model of the world. (It is an interesting question as to why others have not). Other models may maintain symbolic or iconic information, but are not visually detailed throughout. They are often conceptual models (e.g., T. Gao, Gao, Li, Sun, & Shen, 2011). Again, our low-level sensory system is intended to support a capacity limited conscious perceptual system. Other models may *hold* detailed information throughout, but these are not *maintained*. Instead, they must be explicitly erased upon movement to prevent smear (Francis et al., 1994).

A.6.3 First to demonstrate advantage of flowcentric reduction

Our internal world model exists within an active flow of information, i.e., it is *flowcentric*. As such, it has properties that differentiate and often improve upon a *neurocentric* reduction. Firstly, we have demonstrated that the flow is always spatiotopic and hence stable across eye movements – an ideal locus for consciousness. As a consequence, there is no image smear. Neurons are neither retinotopic nor spatiotopic and don't hold such memories, decaying or otherwise – the previous and next chapters elaborate. This may sound absurd, but remember that without eye motion, one cannot tell

the difference between a *flowcentric* or *neurocentric* reduction, nor the difference between retinotopic or spatiotopic frames of reference. One has to ask how a *flowcentric* reduction improves upon a *neurocentric* one.

It is because motion has been considered as fundamental that our model is not only invariant to motion, but actually seems to account for many phenomena by taking advantage of it, e.g., the arthroscopic phenomenon. For example, under a *neurocentric* reduction, neurons can have structurally perfect visual memories, although these would be engineered to decay in ~100ms. However, under eye motion, these contribute to smear (even with decay) so such perfection is wasted – the eyes are almost never at rest, even when supposedly fixated (and even when we dream).

The concept of memory and persistence make less sense under a *flowcentric* regime. Firstly, the structure of "memory" distorts over time – it does not uniformly degrade – however eye motion will refresh the content rather than produce smear. The more the eye moves, the more that is maintained – indeed biological eyes are always in motion, even during fixation It also combines the sensorimotor account that shunned representation with an account that has an internal model, albeit non-representational.

Secondly, there is no time constant to "memory" decay. With some stimulus, interference leads to quick distortion, but with others, memory can remain faithful. Certainly this may complicate analysis, but it can also eliminate the many different proposed memory systems that are defined according to the many different time constants observed.

A.6.4   First to demonstrate a unified model

The demarcation point for how large a set of phenomena can be considered as unified is arbitrary. We certainly hope to show that we have at least a potential for unification. However, we are also the first to demonstrate a robust to stimuli, parameter free and *model free* architecture – no changes of any kind have been made to the ECM architecture.

By model free, we intend to differentiate ourselves from the use of a single architecture by multiple incommensurable models, i.e., these models encompasses an enormously larger set of additional free parameters. We have a single model that happens to have some accidental architecture (Booch, 2006) but with a principled use of an Emergic Network – a realization of a process based epistemology and metaphysics.

Part of the robustness is in avoiding differential equations and using high-level code. Another part is in harnessing the computational power of dendrites – neurons are trivial by example – and this is another move away from *neurocentric* thinking.

A.6.5   Contributions

The *flowcentric* reduction is our most important ontological contribution to empirical cognitive science. However, it is an outgrowth of the emergic approach which is an even greater epistemic contribution. Unfortunately, that must await future research and dissemination.

A.6.6   Future

In the future, the model will be extended to account for object motion as well, but this may require the complications of segmenting visual information across a set of processing *layers* (as distinct from spatiotemporal *levels* which we already have).

Appendix B: An emergic model of filling-in the foveal blue scotoma

At the very center of every human eye lies a small yet significant 0.4° region devoid of blue (S) cones. This foveal blue *scotoma* renders us effectively blue-yellow colour blind – a form of local small field *tritanopia*. Of course, we are never aware of this defect as it is covered-up by the visual system. The supposed filling-in must be extremely fast and precise, as this region is at the very center of overt spatial attention and it is the region of highest spatial acuity – we would notice if something was amiss.

In this chapter, we present a model of filling-in of the foveal blue scotoma based on the Emergic Cognitive Model (ECM) architecture (p323). While ECM does not have any specific mechanism dedicated to filling-in, we hypothesized that both of its basic functions – *handling missing data*, and *maintaining information coherence* – could interact with experimental stimuli and measurements to cause this filling-in phenomena to emerge in a precise and timely manner. As part of characterizing the relative contributions of the two basic functions to the main phenomena, we also discovered that the model would cause border processing to emerge, and we analyze this effect in greater detail – it is a prerequisite for the precision of filling-in.

No other computational models have been proposed to explain this specific phenomenon, so we will contrast our solution to ones that handle the similar phenomenon of filling-in over the optical blind spot, as further investigated in "5   Chapter:" (p161). The major difference between the two is that there are no receptors within the blind spot while there are two colour sensitivity classes within the foveal blue scotoma.

B.1   Introduction

It is quite easy for anyone to characterise some of the functional characteristics of their own optical blind spot or *scotoma*. No special equipment is required. For example, follow the instructions that go along with Illustration 192.

<p style="text-align:center; font-size:3em; font-weight:bold; border:1px solid black; padding:0.5em 1em;">O   X</p>

Illustration 192: Find your blind spot
Close right eye. Position left eye just above X. Focus on X. Slowly move away. The O will disappear at some point ~30cm, and reappear if you continue moving away.

Having found your optical blind spot, you can even move your finger up from beneath the 'O' and "through" your blind spot – the fingertip should be swallowed up by the blind spot until it makes its way to the other side, at which point it will be completed, or filled-in. We will investigate such filling-in of the blind spot in "5   Chapter:" (p161). For now, we note that the 5.5° wide blind spot is quite large and in the periphery compared to 2° wide fovea which consumes almost half of the entire visual processing system (Wandell, Dumoulin, & Brewer, 2007). The blind spot is ordinarily not noticed partly because it is located on opposite sides of each eye and therefore each can cover off (or stereoscopically fill-in) for the other. Even with one eye closed, it is still not ordinarily noticed as monocular filling-in also occurs. One reason not to notice the optical blind spot is that spatial attention is not ordinarily directed to one's blind spot, 15° to the temporal side – it takes a special stimuli such as Illustration 192 to allocate attention there.

What is not commonly known, and much harder to show, is that within the fovea, and within its foveola, there is central 0.4° (25′) region that is devoid of blue (S) cones (Curcio et al., 1991; Magnussen, Spillmann, Stürzel, & Werner, 2001; Williams, MacLeod, & Hayhoe, 1981). This region is critical for high-acuity vision as it contains the smallest diameter red (L) and green (M) cones. Effectively we all suffer small field *tritanopia* – being blue-yellow colour blind in that region.

Illustration 193: Foveal blue-free zone. From (Curcio et al., 1991), Fig. 5. Blue cones distributed heterogeneously but missing in center that could fit a 0.35° disk. Δ indicates highest cone density. S-superior; N-nasal; I-inferior; T-temporal. Bar=100μm.

region. Yet, somehow this defect remains unnoticed. Perhaps defect is too harsh a word.

The lack of awareness is much harder to explain for the blue scotoma than for the optical blind spot. The foveal scotoma is always in the center of overt spatial attention or fixation area (Wald, 1967), and with both eyes looking at the same point in space, one eye cannot cover off for another. Certainly it is small, but it can be made sufficiently salient under experimental conditions involving afterimages where its ragged borders can be detected (Magnussen, Spillmann, Stürzel, & Werner, 2004).

Illustration 194: Foveal blue scotoma. Revealed by method of afterimage formation after blue (450nm) adaptation (Magnussen et al., 2004), Fig. 1.

This is shown in Illustration 194. Flickering monochromatic short-wavelength lights can also reveal the foveal blue scotoma (Magnussen et al., 2001).

Ordinarily, foveal tritanopia reveal its presence indirectly, by a difficulty in discerning blue from green (Willmer & Wright, 1945), yellow from white, or blue from black (Hartridge, 1944) as shown in Illustration 195. Indeed, it was first psychophysically analyzed via colour
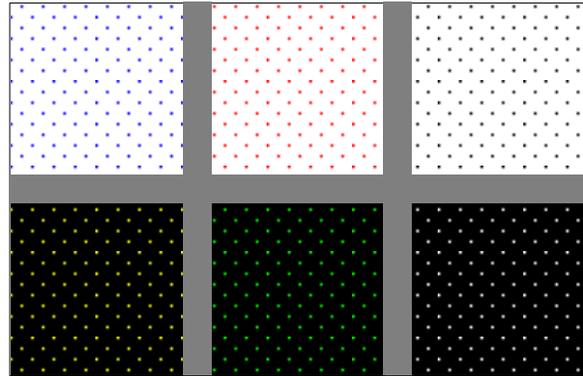
Illustration 195: Difficulty in discerning white from yellow and black from blue

matching experiments using narrow fields of view (König, 1894; Thomson & Wright, 1947) or point sources of light. Discrimination is improved by eye movements (McCree, 1960). Hartridge (1945) also found that tritanopia could occur in the periphery and is also eradicated by motion. We presume that this is due the low density of blue cones – 5-10% in the periphery (Calkins, 2001), and this form of filling-in has been covered in "2 Chapter:" (p9).

Ecologically informed questions may be more relevant. Magnussen et al. (2001) asked why a tiny dark spot was not seen when looking across the blue sky. But, of course, the scotoma area is not blind, just S cone insensitive. Perhaps the following question is more precise. When gazing over a stimulus such as Illustration 196, an actual full field colour blind individual with tritanopia will perceive it as shown in Illustration 197. So why doesn't an individual with ordinary vision not see such colour changes at their foveal blue free region, especially as their gaze crosses a border from one colour band to another? This is exacerbated because their spatial attention is directed at this foveal blue scotoma.

One possible answer is an exceedingly quick filling-in mechanism that respects contours, much like the nearly instantaneous mechanism associated with the optical blind spot (Weil & Rees, 2011).

Illustration 196: Stimulus under ordinary vision. **Illustration 197: Stimulus under tritanopia.**

There will be some differences between blind spot and blue spot filling-in. Not all stimuli illicit filling-in for the optical blind spot, but the missing blue must be added back whenever there are actual stimuli for red and green. Such differences are assumptions on our part based on the fact that the foveal blue scotoma is the area of highest acuity and the focus of overt spatial attention. Psychophysical experiments have been extremely limited due to the difficulty of working on such a small area. Mostly, they reassert that such filling-in is a *real* phenomenon with active neurobiological processes (Magnussen et al., 2001, 2004). Note that when using the paradigm of afterimages, the filling-in times are of the slow variety indicating how long it takes for the afterimage to dissipate; they are not the instantaneous filling-in times associated with maintaining the illusion of a homogenously blue field.

B.1.1   Functional Roles

We have referred to small field tritanopia as a defect to be overcome, but it is always worth considering the possible functional reasons for having a blue free region. Firstly,

blue cones are physically larger, so a lack of blue cones serves to improve acuity where it is most urgent. The human blue-free zone is larger than in monkeys, and human peak cone density is 1.8-3 times greater suggesting adaptation to maximize acuity (Curcio et al., 1991). Of course, it is not clear what it is about the human condition that warrants such an increase. Perhaps it was to detect predators (or track prey) over long distances in open country (C. F. Ross & Kirk, 2007). A second possible reason is to minimize chromatic aberration that is worse at the short wavelengths (Wald, 1967). Accommodating chromatic aberration leads to a final design that is maximally informative (Garrigan et al., 2010).

B.2   Emergic Model

In order to explain the phenomena of this chapter, we have taken the Emergic Cognitive Model (ECM) architecture as summarized on page 5, and made no changes.

ECM has a slow filling-in behaviour that appears like diffusion, but respects borders and is best characterized as interpolative (between borders) rather than extrapolative (beyond border) (p57). Larger sized RFs improve diffusion times, but in these experiments, we verify whether larger RFs can also help to fill-in over the extended regions of space belonging to the scotoma.

ECM also has an instantaneous filling-in behaviour due to the shifting of information from one region of space into another due to incessant eye motion (p171). In these experiments, we verify whether blue information can be shifted into the foveal blue scotoma.

The *handling missing data* function would be invoked in both cases. When the eye is stationary, top-down interpolated blue values will be used. When the eye is moving, lateral shifted blue values will be used. The *maintaining information coherence* function performs the shifting of information to compensate for eye movement. It also requires values to have positional information, and it is this that enables the emergence of border effects. In these experiments, we verify whether borders are indeed respected when filling-in.

B.3   Test 1-3: Full, half and quarter shaped white stimulus

When blue is subtracted from white, then yellow ensues. Thus, when looking at a fully white patch, it should initially appear as yellow at the foveal blue scotoma. In these tests we measure the time for the yellow to fade-out. We could equally call this phenomenon the filling-in of white or the filling-in/addition of the missing blue. While the last might be the most causally accurate, it is descriptively poor because the additional blue is never perceived for a white stimulus – it would be perceived for a dark stimulus that would appear initially black at the foveal blue scotoma.

Leibovitz (p9) already demonstrated the filling-in of blue-free patches with blue in order to perceive a homogeneous blue stimulus as homogeneous. So here we extend the results to a homogeneous white stimulus. Filling-in of blue was especially problematic due to the size of the gaps. Here, the foveal blue scotoma is even larger. As suggested, larger sized RFs could help, and these are tested throughout this chapter.

Stimuli with borders were used in test 2 and 3 to ensure that filling-in respects borders. In test 2 we verify that the blue filling-in behaviour is not excessive and is bounded by one border. Note that while these experiments have been performed for the optical blind spot, they have not been investigated for the foveal blue scotoma. We assume similar behaviour.

Not all stimuli that cover the optical blind spot get filled-in. They must fall across two quadrants of the scotoma, e.g., Illustration 198, with longer lines fully completed and shorter ones only partially (Kawabata, 1983). Periodic stimuli are more effective, but this is considered to involve higher visual areas (Kawabata, 1982; Komatsu et al., 2000). So we will use solid shapes instead.

Another reason to use solid shapes is because these are the simplified black and white stimuli used by alternative computational models, e.g., Illustration 118 on page 152, especially as they provide for distinct border contrast. M. A. Cohen & Grossberg (1984) have an explicit Boundary-Contour System (BCS) used to initiate and limit diffusive filling-in as shown in Illustration 199. We want to ensure that ECM similarly respects borders.

Illustration 198: Half stimulus that is filled-in.
Adapted from (Kawabata, 1982).
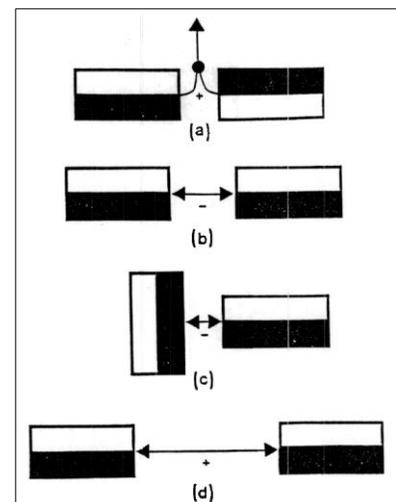
Illustration 199: Explicit boundary-contour RFs.
From (Cohen & Grossberg, 1984), Fig. 1. (a) RFs sensitive to orientation and contrast amounts, but not direction; (b) similar orientations compete over space; (c) different orientations compete over percept; (d) activated/aligned orientations cooperate forming real or illusory contours.

238

In test 3 we verify that the blue filling-in behaviour is not excessive and is bounded by two borders leading to a corner. Note that while these experiments have been performed for the optical blind spot, they have not been investigated for the foveal blue scotoma. In this particular case, there are reasons to assume that the two will not exhibit similar behaviour.

Not all stimuli that cover the optical blind spot get filled-in. They must fall across two quadrants of the scotoma as occurs in Illustration 198 but not in Illustration 200. We assume the filling-in phenomenon under the foveal blue scotoma must occur and this differs from the optical blind spot because in this case, the red and green photoreceptors do provide image content and their erasure or colour change would be salient. Only blue

Illustration 200: Quarter stimulus that is not filled-in for optical blind spot.
Adapted from (Kawabata, 1982).

needs to be filled-in and overlaid on top of the red and green information. Moreover, attention is drawn towards corners (Kinser, 1999) so the absence of filling-in may be too salient.

B.3.1 Methods

The methods used are identical to those of page 36 except as indicated herein.

B.3.1.1 Subject analysis summary

A virtual agent was stochastically generated based on the development parameters detailed in the supplementary material (s309). Its retina is shown in Illustration 201. The same parameters were used to generate a stochastic clone for each of the six experiments in this chapter.
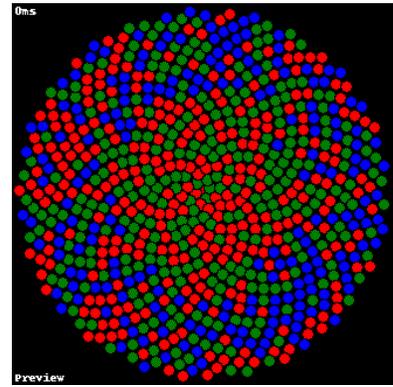
Illustration 201: Retina for foveal blue scotoma

To focus characterization on the filling-in of the blue scotoma phenomenon, a retina with two regions was generated. The inner region was devoid of blue (S) cones having an equal ratio of red (L) and green (M) cones. The outer region had an equal ratio of all three cone sensitivities. The size of the cones did not vary with eccentricity.

The generated retina developed a foveal blue scotoma with diameter of 32.99′ and an outer region with diameter 65.89′. These sizes are relative and not intended for comparison to human data.

B.3.1.2 Stimuli

A full white, half white and quarter white set of images over a black background were presented to the cyclopean agent at time zero. They covered the whole visual field of view for the entirety of the test. The stimuli is fully described in the supplementary material (s310, s311, s312), and are shown in Illustration 202, Illustration 203 and Illustration 204.
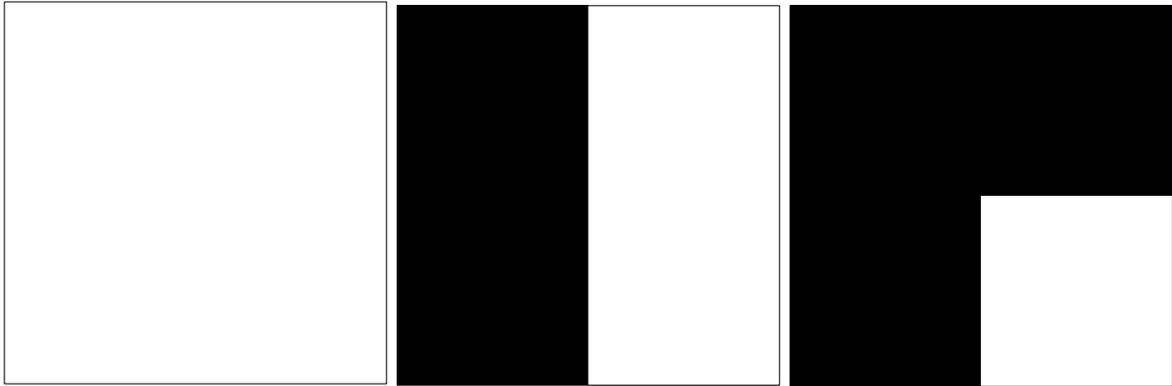
Illustration 202: Test 1 full white stimulus

Illustration 203: Test 2 half white stimulus

Illustration 204: Test 3 quarter white stimulus

B.3.1.3 Procedure

At time=0, the stimulus was presented to the single eye of the virtual agent. The agent was "instructed" to keep its eye stationary (beyond involuntary jitter). The artificial brain was monitored for the first 32 time ticks.

B.3.2 Results

The visual frame-by-frame response results with precise timing information are detailed in the supplementary material (s317, s328, s339). When presented as an animation (Leibovitz, 2012a), all the results for this chapter (and additional stimuli reversing white with blue for each test) can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=165
```

In this section, we summarize and discuss various highlights. We also extract filling-in times.

B.3.2.1 Yellow fading-out behaviour at RF Level 1

For tests 1 thru 3, the dynamic behaviour of yellow fading-out is summarized in Illustration 205, Illustration 206 and Illustration 207. Eye jitter starts at tick=7. While

every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
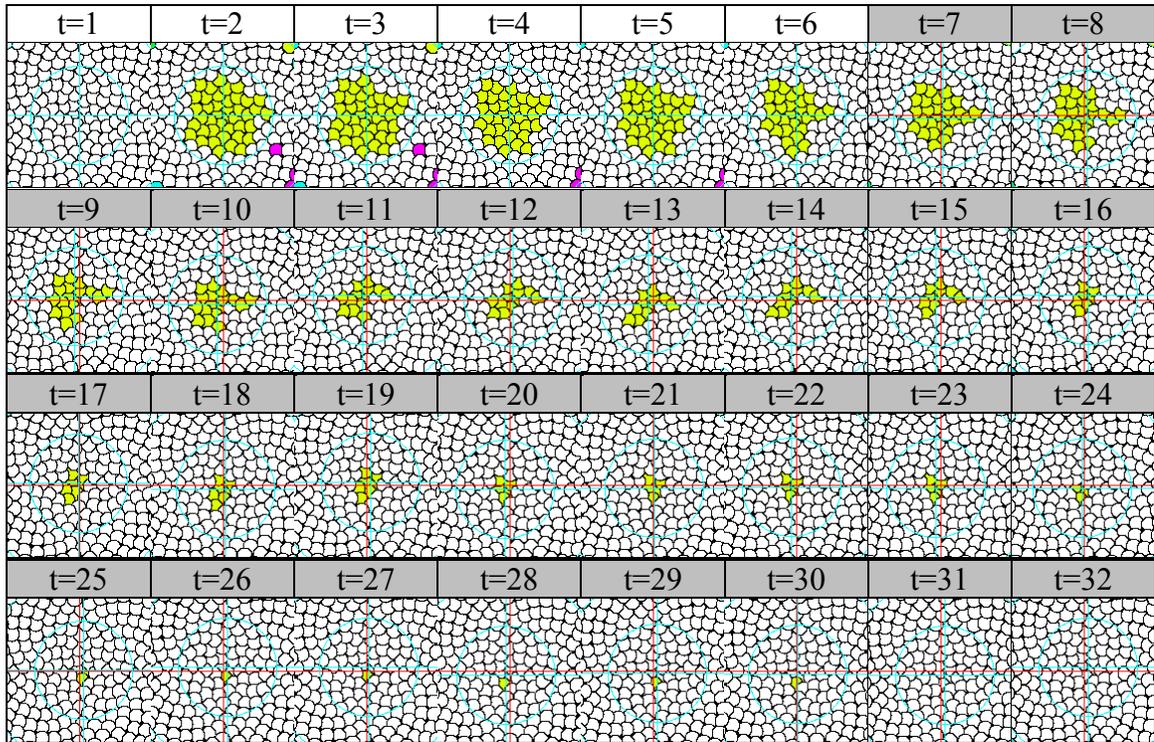
Illustration 205: Test 1 (Full) yellow fading-out behaviour
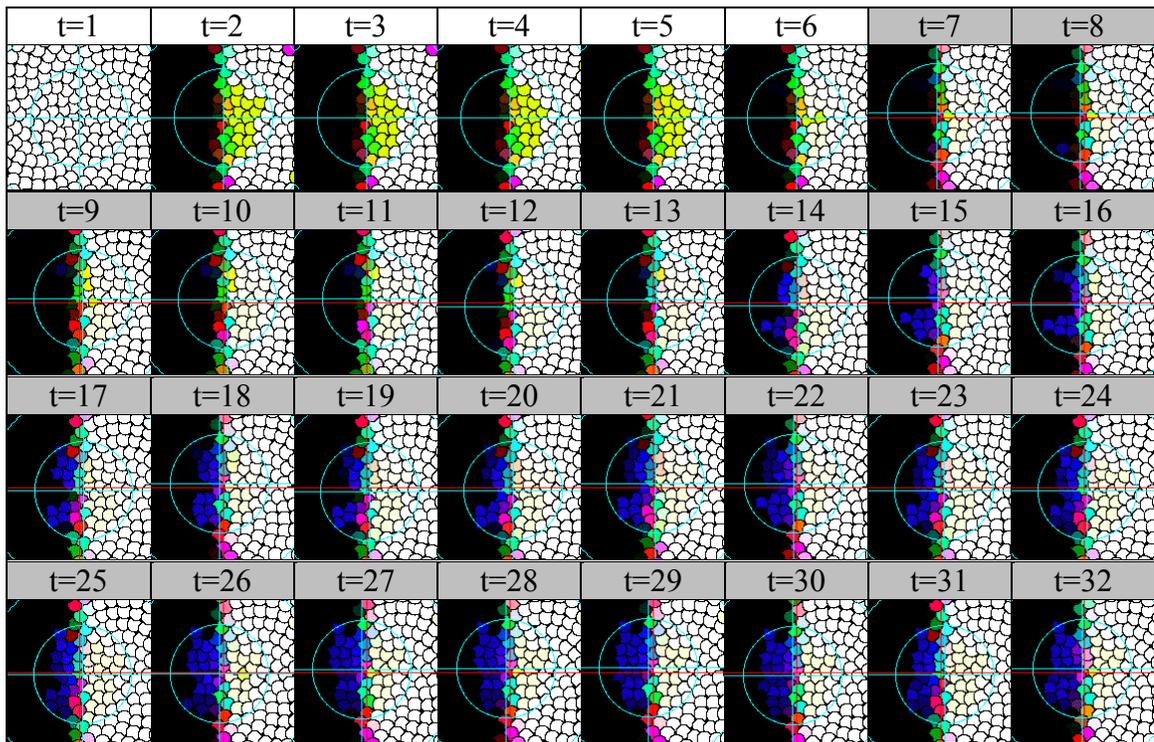
Illustration 206: Test 2 (Half) yellow fading-out behaviour

t=1	t=2	t=3	t=4	t=5	t=6	t=7	t=8
t=9	t=10	t=11	t=12	t=13	t=14	t=15	t=16
t=17	t=18	t=19	t=20	t=21	t=22	t=23	t=24
t=25	t=26	t=27	t=28	t=29	t=30	t=31	t=32
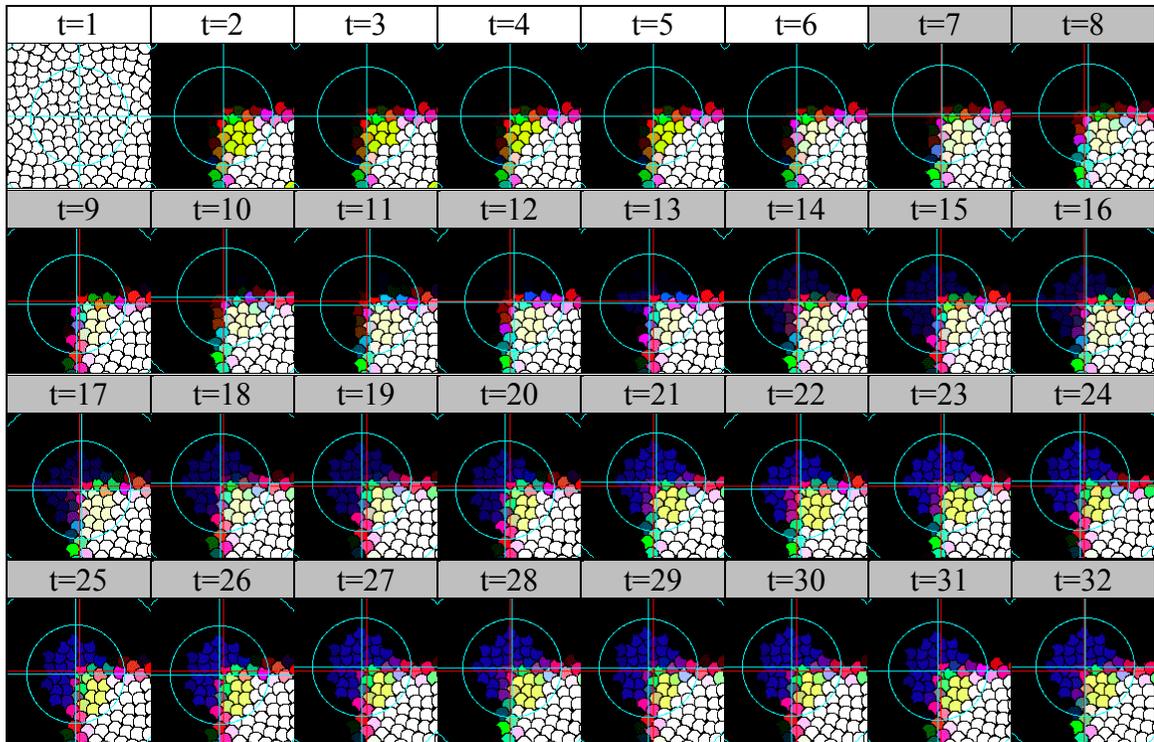
Illustration 207: Test 3 (Quarter) yellow fading-out behaviour

B.3.2.2   Yellow fading-out performance

at RF Level 1

For test 1, the last yellow RF fades out at

tick 31 for a fading-out elapsed time of 29

ticks. Details are shown in Illustration 208.

t=30	t=31
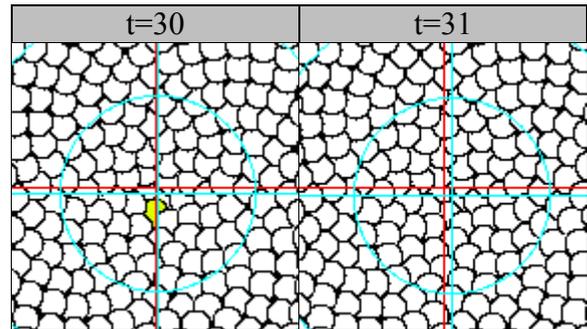
Illustration 208: Test 1 (Full) yellow fading-out performance

For test 2, the last yellow RF fades out at

tick 14 for a fading-out elapsed time of 12

ticks. Details are shown in Illustration 209.

t=13	t=14
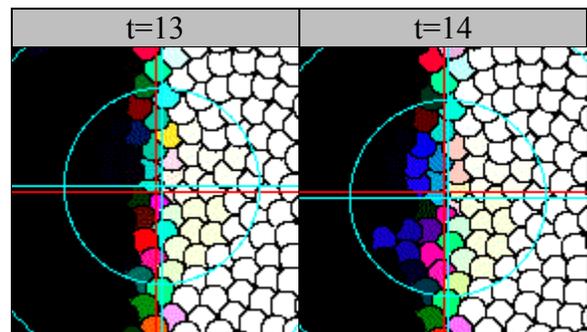
Illustration 209: Test 2 (Half) yellow fading-out performance

For test 3, the last yellow RF fades out at tick 6 for a fading-out elapsed time of 4 ticks. Details are shown in Illustration 210.

B.3.2.3 Border processing

Test 1 has no borders. For test 2, filling-in processes respect borders for the first 13 ticks. For test 3, borders are respected for the first 12 ticks. These are detailed in Illustration 211. Beyond these times, artefacts arise that will be described next.

B.3.2.4 Blue and yellow filling-in artefacts

For test 2, blue starts invading the black region at tick 14 and this is almost completed by tick 32 as shown in Illustration 212.

For test 3, blue starts invading the black region at tick 13 and this is completed by tick 22 as shown in Illustration 213. By this time, parts of the white region have re-yellowed somewhat.
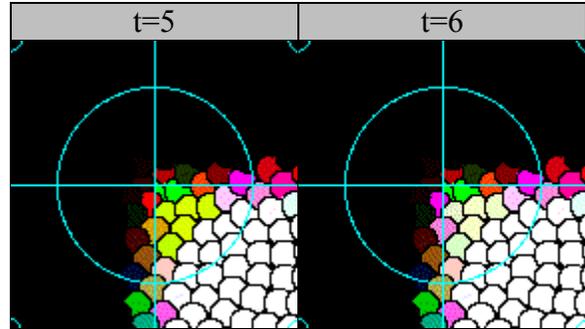
Illustration 210: Test 3 (Quarter) yellow fading-out performance
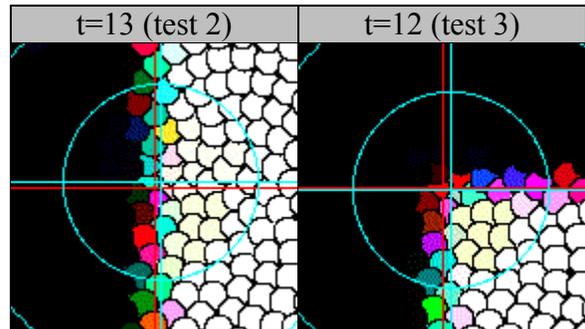
Illustration 211: Test 2-3 border processing
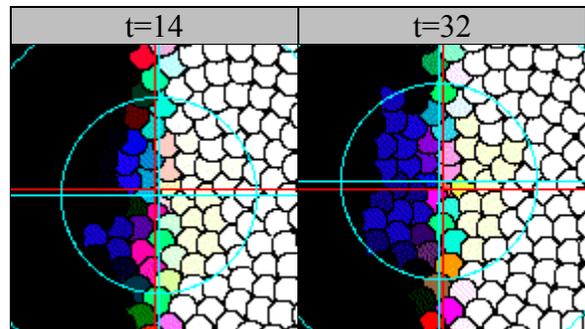
Illustration 212: Test 2 blue filling-in artefact
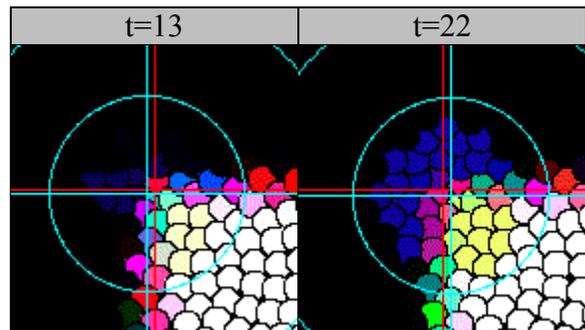
Illustration 213: Test 3 blue and yellow filling-in artefacts

244

B.3.3   Discussion

B.3.3.1   Terminology: Fading-out, filling-in, diffusion

The yellow fading-out description is better explained as blue "filling-in". Even this is incorrect as filling-in implies an active diffusive (most common) or non-diffusive spreading process (Francis & Ericson, 2004) which ECM does not have. Indeed, the behaviour in test 1 indicates that a single yellow RF remains for 5 ticks, and that would not occur under diffusion – it would have lasted at most 1 tick. We retain the unquoted *filling-in* or *fading-out* terminology for descriptive, i.e., epistemic reasons. A causal ECM explanation involves a) summarizing information going upwards, b) using summary statistics reflected downwards while respecting borders and c) giving priority to lateral information that is shifted due to jitter. There are no simple mechanistic reasons nor terms in ECM as all processes interact for emergic effects.

B.3.3.2   Border processing

We have demonstrated that filling-in respects borders. However, ECM has neither explicit filling-in processes nor explicit border processing mechanisms. Our behaviour is emergent due to the interpolating but not extrapolating behaviour (p57). Because all our emergic values have position information, this more complicated encoding mechanism can simplify the need for dedicated neural structures such as the Boundary-Contour Systems (BCS) within Grossberg models (Cohen & Grossberg, 1984). There is a tradeoff between the structure of information representation and the structure of neural systems and functions.

B.3.3.3 Performance

In all cases, the apparent filling-in of blue does complete (29, 12 and 4 ticks). The fading-out of yellow is fastest for the quarter white stimulus. Nevertheless, and without empirical support, we still suspect that this is still too slow and would be noticed at the high-acuity and overtly attended blue scotoma region. In Test 4 we will explore a faster version of fading-out related to eye movement. Another way to improve performance is by having even larger level 2 or 3 RFs at the tritanopic location (p48). However, this might reduce acuity and spread the actual filling-in of blue too far. Tests 2 and 3 verified that the "filling-in" of blue respects borders. While there exists neurobiological evidence for larger RFs at the optical blind spot (Komatsu et al., 2000, 2002), no similar investigation has been performed for the foveal blue scotoma.

B.3.3.4 Blue filling-in artefact

The eventual filling-in of the black region with blue, i.e., not respecting borders, is unexpected. For the optical blind spot which is not ordinarily the focus of attention, this may not be a problem. Indeed blue is hard to discern from black (Hartridge, 1944). For the foveal blue scotoma, we suspect that this would be salient. It could, of course, be a genuine empirical prediction that would need to be verified, but we suspect it is a design artefact.

Nevertheless, there is indeed a case of a similar phenomenon as shown in Illustration 214. McCamy (2003)
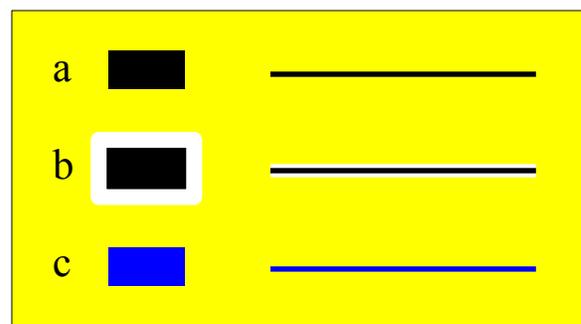
Illustration 214: Blue filling-in of black.
Adapted from (McCamy, 2003), Fig. 1. Line b) is black but appears filled-in with blue.

246

suggested that this phenomenon was related to the mixing of *successive contrast* with *simultaneous contrast* under rapid eye movement over small fields of view. Nevertheless, McCamy was careful to point out that his research did "not establish that connection or any other plausible mechanism." Perhaps we have stumbled upon an inadvertent mechanism? We will not analyze it further but simply reiterate that when a distributed system is defined at a fine-grained level with massive recurrence, many emergic effects will ensue.

B.4   Test 4: Quarter white stimulus with saccade

The previous experiments were not ecologically plausible as the eyes were fixed and only experienced jitter. We hypothesize that under saccadic eye movement, filling-in will be instantaneous because the correct amount of blue would have been shifted in from where the eye could detect blue. In other words, no dedicated yet slower mechanism is required. In this test, we validate our hypothesis.

B.4.1   Methods

The methods used are identical to Test 3 except as detailed herein.

B.4.1.1   Stimuli

The stimulus was similar to that in Test 3, but shifted off to the right, so that the white quadrant is just to the right of the foveal blue scotoma. This is shown in Illustration 215 and is fully described in the supplementary material (s313).
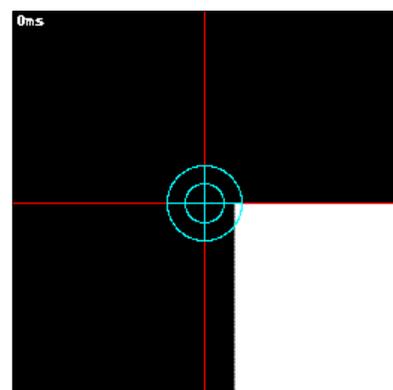
Illustration 215: Quarter white shifted stimulus shown with overlay

B.4.1.2   Procedure

As in Test 3 but with the additional instruction to initiate a continuous saccade to the right by 4′ per tick. Only 16 ticks were monitored as the eye started to move beyond the monitoring space. It is understood that motor planning may incur a delay before actual movement starts.

B.4.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s350). In this section, we summarize and discuss various highlights. We also extract filling-in times.

B.4.2.1   Yellow fading-out behaviour

The dynamic behaviour of yellow fading-out is summarized in Illustration 216. Eye movement (with jitter) starts at tick=7. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
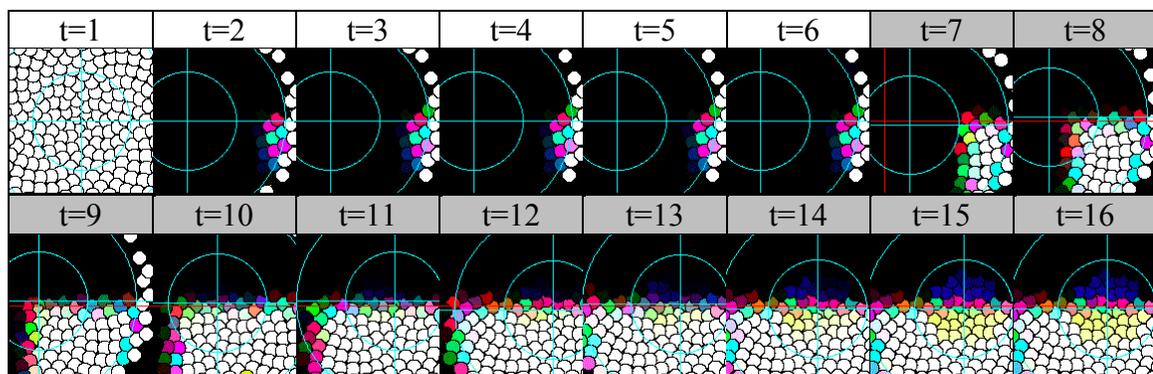
Illustration 216: Yellow fading-out at RF level 1 with quarter white stimulus under motion

B.4.2.2   Yellow fading-out performance at RF Level 1

Yellow is never apparent (and need not
fade out) as the eye saccades over the
white stimulus. Details are shown in
Illustration 217.

t=8	t=9

**Illustration 217: Yellow never perceived under
motion**

B.4.3   Discussion

B.4.3.1   Performance

There is no initial yellow artefact within the blue scotoma that needs to fade-out. In other
words, the fading out is instantaneous taking zero time. This is because the RF at level 1
always has two detailed colour values available upon saccade. It has the bottom-up
information from the photoreceptors (missing blue in the foveal blue scotoma), *and at the
same time*, it has information shifted laterally from the previous region where blue
photoreceptors where available. If blue is not available from below, it can be instantly
utilized from beside as part of the *handle missing data* function of ECM. Our hypothesis
has been confirmed.

Being so fast, perhaps the other solutions to improve performance may not be required.
For example, we do have much larger RFs over the blue scotoma for this chapter as
compared to other chapters. Perhaps standard sized RFs can be used throughout. Such
optimizing questions are for future consideration.

B.4.3.2   Blue and yellow filling-in artefacts

Neither blue invading black nor yellow re-invading white artefacts occur in this test as correct colour information is continuously shifted in from where it is known. So while we consider those artefacts as discussed in the previous experiments as design flaws, we do not worry over much as they do not appear in an ecologically relevant environment.

B.5   Test 5: Arrow stimulus

The typical filling-in of the optical blind spot test involves the use stimuli in the shape of lines or bars, (e.g., Matsumoto & Komatsu, 2005). So in this test we determine how long it takes for a blue line to be filled-in with blue when it sits across the foveal blue scotoma. In this test, the eyes do not move.

B.5.1   Methods

The methods used are identical to Test 1 except as detailed herein.

B.5.1.1   Stimuli

An image of a blue arrow over a grey background was presented to the cyclopean agent at time zero. It covered the entire visual field of view for the entirety of the test. The stimulus is fully described in the supplementary material (s314), and is shown in Illustration 218 where the inner cyan coloured circle demarcates the extent of the foveal blue scotoma – the arrow should appear as black therein.
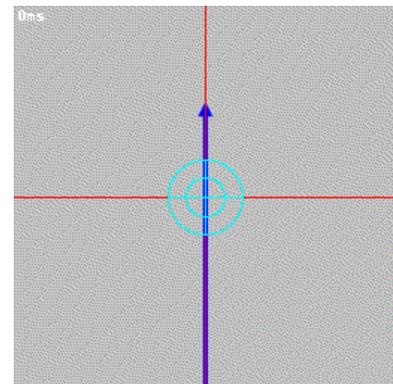
Illustration 218: Blue arrow stimulus
Shown in view with overlay.

B.5.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s356). In this section, we summarize and discuss various highlights. We also extract filling-in times.

B.5.2.1   Filling-in behaviour

The dynamic behaviour of the various filling-in phenomena is summarized in Illustration 219. Eye jitter starts at tick=7. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
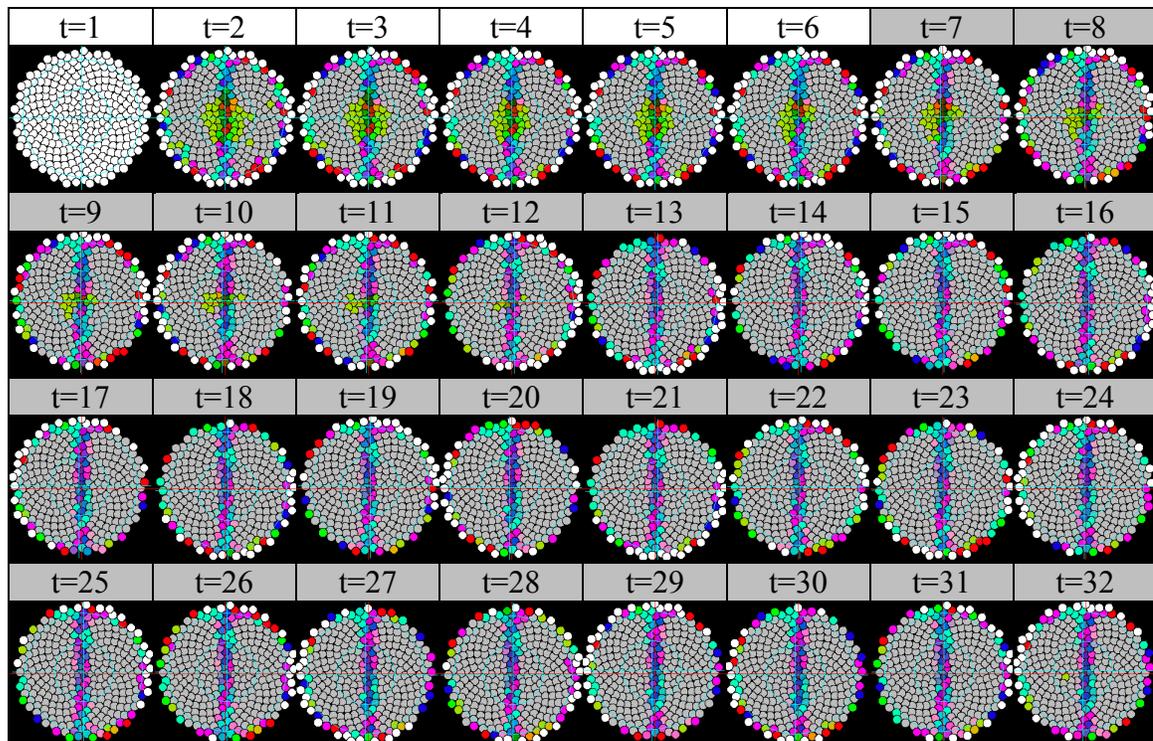
Illustration 219: Filling-in at RF level 1 with arrow stimulus

B.5.2.2 Yellow fading-out performance at RF Level 1

At tick 13, the last yellow RF fades out for a fading-out elapsed time of 11 ticks. Details are shown in Illustration 220.
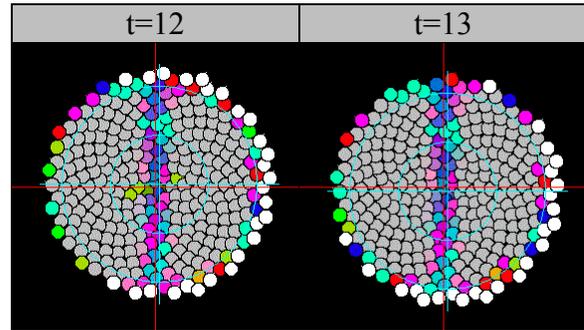
Illustration 220: Yellow fading-out performance

B.5.2.3 Blue filling-in performance at RF Level 1

At tick 12, the center of the foveal blue scotoma does show the arrow as filled-in with blue. Nevertheless, this is not a clearly distinguishable blue throughout the arrow stimulus. At tick 17, the arrow is more clearly defined as blue as shown in Illustration 221.
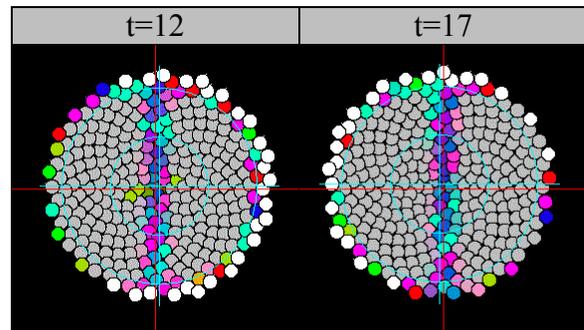
Illustration 221: Blue filling-in performance

B.5.3 Discussion

B.5.3.1 Performance

We suspect that the fading-out of yellow and filling-in of blue were likely too slow and would be salient for the foveal blue scotoma. A better solution was implied by Test 4, and we will verify it on a point stimulus in Test 5.

B.6 Test 6: Point stimulus

The foveal blue scotoma makes it hard to distinguish white from yellow, and black from blue (Hartridge, 1944) as shown in Illustration 195 on page 234. This is especially true under a colour matching paradigm using narrow fields of view (König, 1894; Thomson & Wright, 1947) or point sources of light. Nevertheless, discrimination is improved by eye movement (McCree, 1960).

In this test we verify 1) whether eye motion improves filling-in of a blue point source stimulus, and 2) how quickly such filling-in occurs. In essence, it is a combination of Test 4 with Test 5.

B.6.1 Methods

All the methods are identical to Test 4 except as indicated herein.

B.6.1.1 Stimuli

An image of a blue point over a grey background was presented to the cyclopean agent at time zero. It covered the entire visual field of view for the entirety of the test. The stimulus is fully described in the supplementary material (s315), and is shown in Illustration 222 where the inner cyan coloured circle demarcates the extent of the foveal blue scotoma.

Illustration 222: Point stimulus
Shown in view with overlay.

B.6.1.2 Procedure

After the point source stimuli appears, the agent was instructed to saccade its one eye to the right by a continuous 4′ per tick. Results were record for 24 ticks.

B.6.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s367). In this section, we summarize and discuss various highlights. We also extract filling-in times.

B.6.2.1   Filling-in behaviour

The dynamic behaviour of the various filling-in phenomena is summarized in Illustration 223.
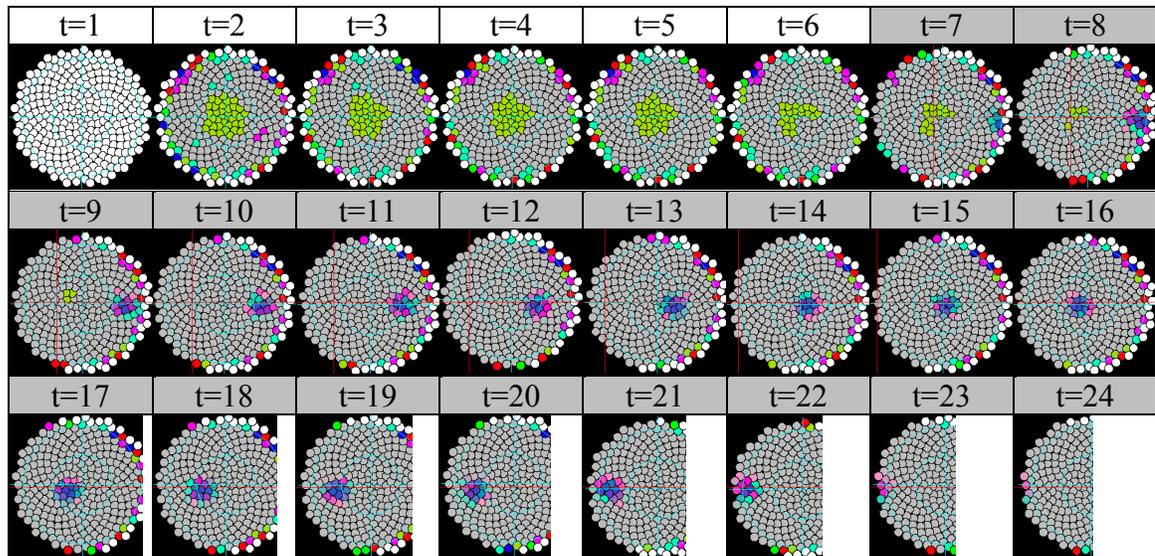
Illustration 223: Filling-in at RF level 1 with point stimulus

Note that at tick=17, the retina starts to move beyond the area preallocated to show a stable headcentric display. Thus, the image truncation.

B.6.2.2   Yellow fading-out performance at RF Level 1

At tick 10, the last yellow RF fades out for
a fading-out elapsed time of 8 ticks.
Details are shown in Illustration 224.
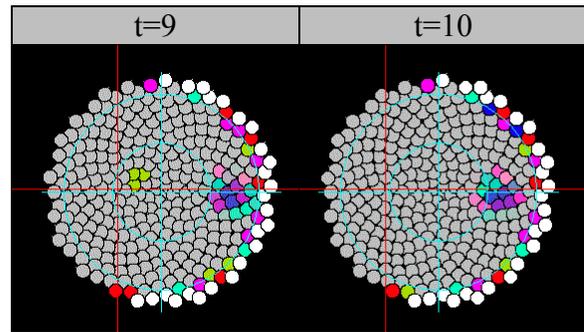
Illustration 224: Yellow fading-out performance

B.6.2.3   Blue filling-in performance at RF Level 1

At tick 10 (as shown above), the point
stimulus first enters the foveal blue
scotoma and it comes in fully formed in
blue. The effective filling-in of blue times
is instantaneous as shown further in
Illustration 225.
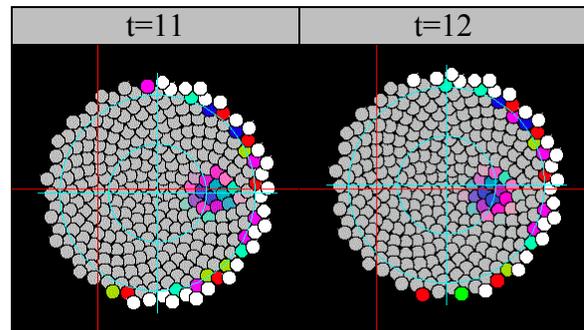
Illustration 225: Blue filling-in performance

B.6.3   Discussion

B.6.3.1   Performance

If one assumes that a blue stimulus ought to appear as black within the foveal blue
scotoma, then the apparent filling-in of blue is instantaneous for exactly the same reason
as in Test 4. Thus eye movement does help discriminate blue points from black.

B.6.3.2 Realistic Peripheral S-Cones

For simplicity, the ratio of S cones in the periphery was made identical to L and M cones – ⅓ each. In actuality, only 10% of the periphery is made of S cones, leading to local patches of tritanopia (p15). By manipulating the spatial frequency of blue or yellow stimuli, they can be made to disappear in a large field tritanopia phenomenon (Logvinenko & Hutchinson, 2006; Logvinenko, 2001). The implication to our model is that if a thin blue line in the periphery appears as black, then when shifted into the foveal blue scotoma, it will also appear black as no blue information has shifted in. This is desirable as the colour of an object does not change as it crosses the foveal blue scotoma. In any case, Logvinenko makes the case that transient tritanopia implies that the "hypothetical process of filling-in, is not correct" (Logvinenko, 2001). We couldn't agree more – we have no explicit filling-in mechanism, and the ecological case with eye movement is due to shifting.

B.7   General results and discussion

Receptive fields that cover the optical blind spot have been found (Komatsu et al., 2002) and this suggests some active filling-in process. However, there has not been a similar investigation for the foveal blue scotoma. Part of the problem is that this region is very small and within an already highly dense foveal confluence. In this chapter, we investigated the filling-in (and sometimes fading-out) behaviour across the foveal blue scotoma within the ECM architecture, and contrasted with known or assumed human behaviour.

B.7.1 Performance comparison

Psychophysical filling-in times have not been measured for the foveal blue scotoma. We will assume they are at least as fast as filling-in signals elsewhere. Paradiso & Nakayama (1991) found that brightness induction can travel as fast as 6.6-9.2ms/deg. Assuming linearity to our small scale, this translates to an upper bound of 4ms for a blue scotoma of 0.4° – roughly the operating period of a single neuron, and less than a single ECM tick. Blakeslee & McCourt (2008) found induction to be nearly instantaneous suggesting that an active filling-in process from the edge would be too slow. Evidence supports a spatial filtering account. Under eye motion, we offer our own instantaneous model as an alternative at the *sensory* level, although it is not a model of brightness induction, colour constancy, or any *perceptual* phenomenon.

B.7.2 Relation to ECM functions

When the eyes are perfectly stabilized (barring jitter), the <u>slow</u> version of filling-in occurs due mostly to the *handle missing data* function combined with information flows summarizing upwards and filling-in downwards. This provides the illusion of filling-in by a diffusive process. However, "diffusion" is nothing more than an interpolation between summary statistics that retain virtual borders. There is no extrapolation beyond borders. This provides for the precision of filling-in, so that a pie shaped stimulus does not fill-in the entire foveal blue scotoma, but retains its wedge shape.

With slight eye motion, the *maintain information coherence* function ensures that blue values are shifted from actual sensed locations into the appropriate place within the foveal blue scotoma. This maintains perfect precision, and only requires lateral

information flows. Ultimately, the *handle missing data* function is still required to make use of lateral information when available. This version is infinitely <u>fast</u>, or incurs zero delay, as an RF can receive information from a photoreceptor underneath, but if not available, it can use the lateral information. The reinstatement of normal trichromatic visions under slight eye motion is a known physiological behaviour of filling-in (Hartridge, 1945) and is demonstrated in ECM.

B.7.3   Comparison to alternative computational models

There are no computational cognitive models that have been specifically created for the filling-in of the foveal blue scotoma. Indeed, there is very little phenomenology about this topic, and its seems to have been exhausted, e.g., (McCamy, 2003; Tuck & Long, 1990). Instead, we will review computational models for filling-in of the optical blind spot, or other scotomas. There can be no direct comparison because in full scotomas, no red, green nor blue spatial information is available and both contours and colours need to be completed, whereas in the blue scotoma, at best, all the red and green defined contours are present and it is simply a matter of interpolating the blue values. At worst, a blue-only shape would still need to have contours completed. Similarly, scotomas that develop in later life are not perfectly accommodated by neural plasticity, and their filling-in is not as good as for the blind spot or the foveal blue scotoma.

B.7.3.1 Image processing

McManus, Ullman, & Gilbert (2008)
presents a computational model to explain
the filling-in of scotomas that can develop
in later life, e.g., due to retinal
degeneration. The results of their
simulation appear in Illustration 226.
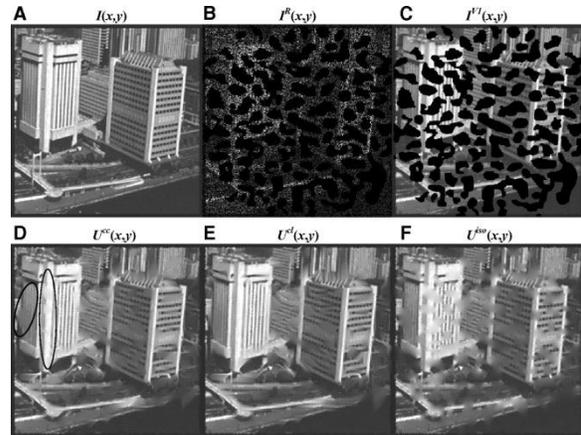
Illustration 226: Filling-in simulation.
From (McManus et al., 2008) Fig.4, showing various stages of visual processing.

We call this an *image processing* account as it uses a standard bottom-up filter based methodology over static pictures (Graham, 2011; Kingdom, 2011) without consideration of processing times. In this case, a further simplification is made in the use of greyscale pictures, and it is not clear whether the proposed algorithms translate to colour processing. For example there are many pathways throughout the visual system. Perhaps this model emulates the achromatic (greyscale) pathway, but then how would they go about filling-in the chromatic pathway? In particular, the luminance pathway is a result of R+G and so the filling-in of blue may require alternative mechanisms. The real problem is that neurobiological evidence is taken ad-hoc without regard to whether it applies to the correct pathway or not. Human vision is colour, not greyscale.

The purpose of this model is less to explain a good filling-in mechanism (which is based on local association fields or AFs (D. J. Field, Hayes, & Hess, 1993)), and more on explaining the visual artefacts due to some neural reorganization involving the permanent shifting of receptive fields. The authors claim that their model "can engender filling-in

that is both robust and consistent with psychophysical reports of perceptual completion", but no empirical comparisons have been. Quantitative comparisons to patients would have been ideal, but even qualitative recreations of their visual deficits would have made convincing evidence. Descriptions of effects on active perception would also have been useful. Moreover, the authors note that the shape of the AFs makes a difference to what kind of scenes are filled in "well". Thus their findings are not robust. Additionally, their use of aesthetically "well" belies an image processing heritage. We would like a good match to human errors, not to image aesthetics. Furthermore, they do hypothesize a top-down contextual modulation to improve performance, but then with any unimplemented hypotheses, these could also provide for perfect vision which is not phenomenally accurate.

Their results do support their claim but it is week evidence. It may well be the case that many compensatory deviations from optimal processing would produce similarly poor results so *any* change of plasticity might do, not just simply theirs. Indeed, a change in the colour pathway may be significant. For example, a similar damage to our model, without any compensation mechanisms, could give similar poor results. This is demonstrated in a latter chapter of filling-in over the blind spot – the slower the blind spot moves around an image, the worse the distortion because dynamic filling-in is instantaneous but filling-in under static eyes takes time (p40).

As a final comment, this particular brand of filling-in will not do for either the optical blind spot or the foveal blue scotoma. For one, it intentionally has too many artefacts, and

secondly, the authors also mention that phenomena such as contour completion are of the wrong kind. Across a blind spot, a line appears *completed*, across a region of empty space, it may appear *illusory*, but not actually present – one cannot use the same mechanism for both.

There are alternative contour mechanisms that unify both illusory and occluded contours (Kalar et al., 2010), but they have not been integrated with the filling-in of surfaces as required within the foveal blue scotoma, nor the optical blind spot. Moreover, the space around the optical blind spot (and we assume the foveal blue scotoma) does differ from normal empty space in that 1) illusory contours do not get completed (Maertens & Pollmann, 2007), and occluded contours may be partially completed (Abadi et al., 2011). A 10° diameter circle falling partially over the blind spot will not have its arc completed (Ramachandran, 1992b). Entire dots may be interpolated even if they do not touch the scotoma boundaries (R. J. Brown & Thurmond, 1993), concentric circles would fill-in (Kawabata, 1982).

When two bars are of opposite colour are situated on opposite side of the blind spot, they are still are completed, but their meeting boundary is not available to conscious awareness with subjects reporting "It's funny; I can tell they meet but I can't say where!" (R. J. Brown & Thurmond, 1993).

B.8   Conclusion

It was once questioned whether the filling-in of the optical blind spot was an illusion (Dennett, 1991) – as it is out of the way and it was possible that perceptual awareness

(attention) was simply not allocated in that area. Of course, now we know that there are some active processes and that these are influenced by attention (Lou & Chen, 2003). In the case of the foveal blue scotoma which is at the very center of vision, it does not seem possible that attention is not allocated in that area while awareness is blind to the stimulus dimension of blue – we can perceive both blue and white stimuli very well. Well, not actually that well. White and yellow are hard to distinguish as are blue and black (Hartridge, 1944) under low light or narrow fields. Perhaps this is the only phenomenal evidence for the anatomical tradeoffs in vision. Nevertheless, under normal vision, one does sense blue very well and one is blind to their foveal blue scotoma – few have ever heard of it.

While filling-in of missing blue information within the foveal blue scotoma may operate similarly to the optical blind spot. There may be differences. For example, a quarter field stimulus will not be completed over the optical blind spot while it would over the foveal blue scotoma as the red and green components of the stimulus are still visible.

In this chapter we presented a computational model of filling-in over the foveal blue scotoma based on the ECM architecture. It meets several performance and phenomenal requirements based on shifting information to compensate for eye movement.

Appendix C: An emergic model of filling-in during blinks

The eyes undergo incessant blinking, yet the world appears stable and undisturbed. The momentary darkness is supressed or masked while visual continuity is maintained, e.g., by temporal filling-in while the eye is closed. There are many similarities between blink suppression and saccadic suppression (Ridder III & Tomlinson, 1995, 1997), for example, change detection is equally diminished over static scenes. However, when viewing *dynamic* scenes, change detection under blinks outperforms those under saccades (Dornhoefer, Unema, & Velichkovsky, 2002). Therefore, the need to account for both eye and scene motion is vital in computational models involving blinks. Moreover, blinks are often timed to occur during gaze shifts when motion is ever present (Evinger et al., 1994; Fogarty & Stern, 1989; P. Fu, Chang, O'Beirne, & King, 2008). Thus temporal filling-in alone may not be sufficient to explain the spatial stability that is also maintained.

Modelling itself is equally vital. When Volkmann (1986) reviewed the previous 25 years of visual suppression research he noted "that most investigators are moving away from the type of dichotomous thinking that attempts to account for all of the findings with one mechanism." His hope for the next 25 years was that "the precise relations among the mechanisms and the different levels of analyses of the systems will be illuminated much further." Today, the one methodology that can specify and control the context dependent use of multiple mechanisms is computational modeling. Surprisingly, the blink research paradigms still in use today do not admit for such models – they retain single theoretical and mechanistic approaches. Research today has worsened matters by multiplying

possible mechanism and phenomenon. As Krekelberg (2010) summarizes for the similar field of saccadic suppression:

> *"These are only two of the many examples of neural evidence that are consistent with one of the possible mechanisms. This embarrassment of solutions suggests that the question "where does suppression take place?" may be ill-posed. Instead, it may be more fruitful to think about the functional role that each area — or even each neuron — plays, and what that neuron would require to function appropriately in the presence of eye movements."*

In this chapter we present a low-level sensory model of filling-in during the blink suppression period and while the eye is in motion. The model is based on the Emergic Cognitive Model (ECM) architecture (p323). ECM can clarify conceptions precisely because it does not prematurely reify them. For example, ECM does not have any specific filling-in mechanism nor any visual memory storage facility to supply the filled-in information. However, we hypothesize that both of its basic functions – *handling missing data*, and *maintaining information coherence* – could interact with experimental stimuli and measurement to cause outward phenomena such as image continuity to emerge, as well as causing the emergence of inward constructs, behaviour, mechanisms and systems such as filling-in and working memory. In other words, rather than model the phenomena directly, we model other functional aspects explicitly, and the phenomena in question emerge indirectly. As always, our results are demonstrated visually and in colour.

C.1 Introduction

In this section we introduce the phenomenology associated with eye blinks – their dynamics, their suppression, and their filling-in. As well, we also introduce the various methodological and theoretical approaches used.

C.1.1 Eyeblink dynamics

On average, the human eye blinks every 4 seconds for a duration or 250-450ms (Volkmann, Riggs, Ellicott, & Moore, 1982). However, this blink rate is extremely task (Karson et al., 1981; Ponder & Kennedy, 1927) and cognitive (Holland & Tarlow, 1975; Shultz, Klin, & Jones, 2011) dependent which will cause trouble to fMRI approaches as we shall shortly see.

Illustration 227: Sample eyeblink dynamics.
From (Volkmann et al., 1980), Fig. 1: (A) Physical eyelid position with dashed lines representing pupil margins; (B) Corresponding electromyographic (EBG) signal.

Volkmann et al. (1980) measured the dynamics of many eyeblinks, one of which is shown in Illustration 227. In this instance, it took 40ms from blink onset to completely cover the pupil. The pupil remained covered for 110ms, and the total blink duration was 350ms. The blink itself can be very fast reaching speeds up to 2400°/s, and it only takes 4ms to cover the 4mm of the pupil (Gawne & Martin, 2000).

C.1.2 Blink suppression dynamics

Volkmann et al. (1980) also wanted to study the temporal

dynamics of *neural* blink suppression which was thought

to be similar to saccadic suppression. However, this

necessitated delivering visual stimuli that was immune to

physical eyelid interference. As shown in Illustration 228,

they shone the light stimulus through optical fibers into

the mouth bypassing the eyelids completely!
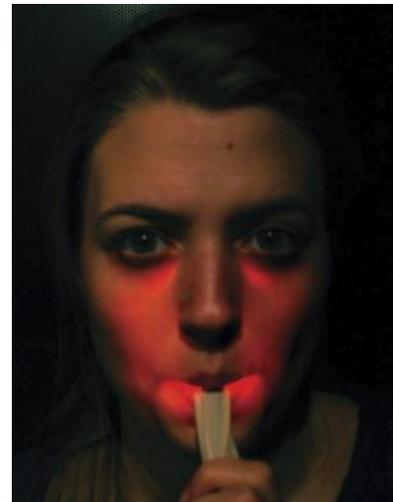
Illustration 228: Eyelid bypass technique of Volkmann et al. (1980).
Image from (Burr, 2005).

Their psychophysical results are shown in

Illustration 229. Of note, suppression started

before blink onset, reached maximum before the

eyelid covered the pupil, and recovered about

200ms after blink onset before the eyes where

fully opened. Such advance suppression is similar

to the suppression that occurs during voluntary

saccades and suggest an extra-retinal signal such

Illustration 229: Visual sensitivity of stimulus time in relation to blink onset.
From (Volkmann et al., 1980), Fig. 2: curves for two participants.

as the efference discharge motor command to close the eye (Volkmann et al., 1982).

Presumably, during this ~300ms of blink suppression, the stimulus would need to be

temporally filled-in.

C.1.2.1 Suppression vs. masking

Because suppression occurs before the eyelid is closed, and released before it is opened,

Volkmann et al. (1982) found it "difficult to imagine how such a finding could be

explained adequately by models of masking." At that time, and today still, it is thought that the mechanisms for blink suppression and filling-in are similar to those of saccadic suppression and filling-in (Burr, 2005). However, today it is recognized that masking is the better explanation for saccades, and *suppression* is a bad term (p69). For example, the insertion of a short blank period after a saccade can eliminate saccadic suppression (Deubel, Bridgeman, & Schneider, 2004; Deubel, Schneider, & Bridgeman, 1996), and this *blanking effect* can equally eliminate blink suppression as well (Higgins, Irwin, Wang, & Thomas, 2009) – it can actually *enhance* displacement detection (Deubel et al., 1996). Surprisingly blinks themselves do not eliminate suppression as do blanks, i.e., blinks are less like blanks and more like saccades. Their extraretinal source of suppression seems to be similar (Ridder III & Tomlinson, 1997), as is their neural ramifications, but not in the early visual areas (Burr, 2005).

The key point is that the undoing of suppression is triggered after the supposed suppression has already occurred, implying that it could not actually have occurred. Information was indeed transmitted, but was masked or unutilized for other reasons. Such conditional masking or "suppression" effects cry out for a computational model of when information is used and why.

Such widespread variations or context sensitivity in visual processing related to blinks may lead to methodological errors and can undermine the results of fMRI protocols (Hupé, Bordier, & Dojat, 2012). For example, (Bristow, Frith, & Rees, 2005; Bristow, Haynes, Sylvester, Frith, & Rees, 2005) examined the effects of *voluntary* blinking vs.

non-blinking ("fixation") over stationary stimuli of two kinds, a simple fixation point vs. fixation point on a checkerboard pattern. Additionally, the amount of illumination was either held constant or was flickered. They found 21 brain regions with reduced BOLD activity during voluntary blinking, clearly
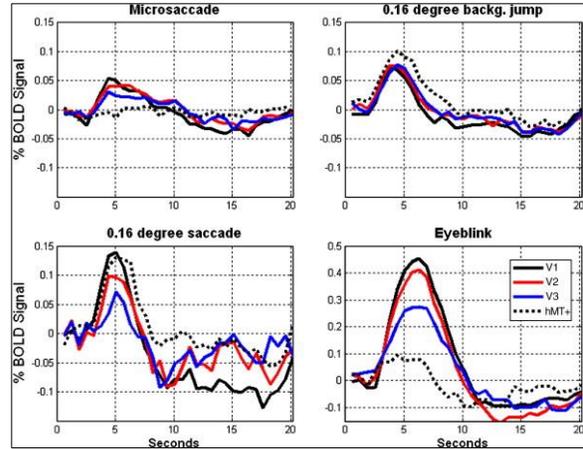
Illustration 230: BOLD signal after microsaccade, image motion, eye motion and eye blink events. From Tse, Baumgartner, & Greenlee (2010), Fig. 7

demonstrating blink suppression. However, Tse, Baumgartner, & Greenlee (2010) found that when using a more ecological relevant paradigm involving eye or stimuli motion, that *spontaneous* blinks were associated with a robust increase in BOLD signals as shown in Illustration 230. Hupé et al. (2012) contend that because the BOLD response is so *strong*, widespread and has great individual variation, that it becomes next to impossible to have a cognitively balanced protocol, and that blinks are therefore, much more likely to trigger correlational artifacts – they are a source of noise to be ignored. In essence, fMRI results average over cognitive processes, and where these cannot be kept identical or perfectly balanced, the results become dubious.

Even if blinks did not contribute noise, any of the increased BOLD responses of Illustration 230 could be attributed to filling-in processing, and arbitrarily of course, any decreases to suppression – as if these are the only two processing choices possible in visual perception. Thus computational models are required to indicate specific functional processes and traditional fMRI results are shown again to average over computational

processing. Newer fMRI approaches are available, but they require an a-priori model selection: "temporal precedence – and causality concepts based on it – are essential in dynamic models of brain connectivity" (Roebroeck, Formisano, & Goebel, 2009). They also have significant limitations (Daunizeau, David, & Stephan, 2009; Valdes-Sosa, Roebroeck, Daunizeau, & Friston, 2011). There is a seeming paradox here. One can use *multiple* fMRI results to induce a functional model, yet that model is used to interpret the fMRI results in the first place. The way out is hinted by italics – a holism of analysis (beyond just fMRI) can be used to induce a model that serves to interpret a single fMRI result. It also implies that any interpretation today is likely to require reinterpretation tomorrow – fMRI results alone cannot be taken too seriously as it is the development of better models that improve the understanding of neuroimaging data (Friston, 2009; Sui, Adali, Yu, Chen, & Calhoun, 2012). Moreover, it is not clear whether the *standard* cause/effect models are appropriate to explain brain dynamics that occur due to *natural* subsystem interactions (Chicharro & Ledberg, 2012).

C.1.3   Temporal Filling-in (Memory) Dynamics

For approximately 44 cumulative minutes/day, while information is masked or suppressed, and certainly while the eyes are closed (Shultz et al., 2011), information must be temporally filled-in across the blink duration to maintain visual continuity. Suppression alone cannot suffice for perceptual stability (Burr & Morrone, 2011). In the case of *saccadic* suppression, the transient
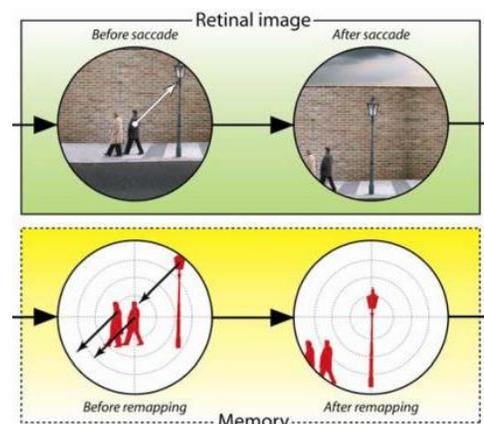
Illustration 231: Saccadic remapping example.
From (Bays & Husain, 2007), Fig. 1: White arrow indicates planned eye movement while black arrows indicate remapping in some kind of memory store.

predictive neural remapping as shown in Illustration 231 could ostensibly fill-in from the external world into a stable spatiotopic homunculus. However, in the case of *blinks*, the external world is not available to be remapped. Some form of memory of the world is required for temporal filling-in. As blinks often occur during gaze shifts (Evinger et al., 1994; Fogarty & Stern, 1989), the filling-in must be spatiotemporal (Schill & Zetzsche, 1995). At this time we are not aware of models, conceptual or computational, that incorporate both temporal and spatial aspects. Indeed there are few models that consider blinks at all.

We will first review dynamics of memory conceptions useful for blink filling-in with fixed eyes then for the more plausible dynamics eyes.

C.1.3.1   Static Eyes & Neurocentric Memory

Most memory models presume that the eye does not move, and that memory for items in outer (visual) space is associated with structures in inner (neural) space. For example, Billock (1997) suggested reverberations between the LGN and V1 as a form of visual short term memory (VSTM) that could be used for temporal filling-in across blinks and saccades. Recurrence is a recognized mechanism to explain the phenomenon of memory (Maex & Steuber, 2009), however feedback need not be *vertical*

Illustration 232: Mean firing rates over time for eyeblinks or stimulus blanking.
From (Gawne & Martin, 2002).

between layers of visual processing, but can also be *lateral* within the same layer, (e.g., Iijima et al., 1996). One of the problems with vertical interactions is that higher layers

have larger spatial magnification factors and there would be an effective loss of fidelity upon feedback. For example, V1 feedback corresponds to an area four times larger than an RF center in the LGN, and slightly larger than the surround (Angelucci & Sainsbury, 2006). The current roles for such vertical feedback are thought to be of modulatory nature, and specifically to sharpen the RFs of LGN neurons and enhance their transmission of relayed signals (Briggs & Usrey, 2011). Indeed, if there is any temporal filling-in across blinks (or darkening) beyond 150ms it does not involve neural activity in V1, V2, V3V/VP nor V4V (Gawne & Martin, 2002) as shown in Illustration 232.

Research into perceptual stability across blinks or saccades take the existence of some form of memory for granted. What is never mentioned in blink research, for example, are the dynamics of afterimage research, even though both might tap into the same underlying systems associated with visual persistence. Similarly, most afterimage research does not look at afterimages formed during voluntary eye closures. For example, visible persistence lasts for 100ms after the stimuli is removed (Di Lollo, Clark, & Hogben, 1988), but this may not be enough time to cover off the entire 200-300ms blink period. What is worse, is that research into visible persistence has shown an inverse duration effect – the longer the stimulus is presented, the shorter the persistence (Coltheart, 1980) invalidating the sensory storage hypothesis (Di Lollo et al., 1994) implicit in conceptual models of blink filling-in, and in the existence of a spatiotopic representation. Finally, we should point out that much of the research conflates persistence and afterimages. Technically, the persistence is available the moment of

stimulus offset (for 100ms) and is likely of cortical origin, while the afterimage is delayed by about 1s and is likely of retinal origin (Di Lollo et al., 1988).

Other problems occur with the concept of *iconic memory* that could provide the information needed to fill-in across blinks. Iconic memory is an outgrowth of the partial report paradigm (Sperling, 1960) and seems to suggest that *information* (possibly categorical) persists approximately 100ms after stimulus offset – not long enough. However, other forms of "iconic" memory can last 250ms (Haber, 1983) or even 1000ms (Graziano & Sigman, 2008). It is not, therefore, a unitary phenomenon and may simply reflect arbitrary discernible phases of information processing (Di Lollo, 1980; di Lollo, 1983). While the locus for iconic memory is cortical and possibly distributed (Keysers et al., 2005) – visual persistence has been evidenced at all levels of visual processing (Haber, 1983) – it is not suitable for filling-in across blinks as it is disrupted by blink processing (Thomas & Irwin, 2006). Haber (1983) has called for the elimination of the iconic memory construct as originally defined to contain static snapshots of a scene, and argued that it could not serve any ecologically relevant function. Nevertheless, it is still studied as the first of three information processing stages forming visual short term memory (Sligte, Vandenbroucke, Scholte, & Lamme, 2010; Zylberberg, Dehaene, Mindlin, & Sigman, 2009).

C.1.3.2 Dynamic Eyes & Emergic Flowcentric Memory

Most models are for static scenes. Gibson (1950) introduced the concept of *optical flow* to account for a more ecologically relevant environment where the agent and visual objects both undergo continuous motion. As shown in Illustration 233, this flow field is information rich and could allow for *direct perception* without the need for synthesis or inferencing. Gibson sought a "psychophysical bottom-up approach to perception" and "scrupulously avoided mention of internal representation"

Illustration 233: Optic flow field superimposed on visual scene. From (Chessa et al., 2009), Fig 3d: Pedestrian demonstrates object motion, while rest of image demonstrates ego motion – the closest part of the scene (bottom of road) has the largest apparent motion vectors.

(Nakayama, 1994) or any intermediate form of memory (Norman, 2001). Today, that approach is part of the standard visual research methodology, e.g., the hierarchy of multiple pattern recognizers (Graham, 2011); now augmented with additional motion related information across the visual field as provided by the numerous kinds of motion sensitive neurons modeled as spatiotemporal filters (Burr & Thompson, 2011). It is, therefore, highly ironic that the concept of optic flow with direct perception was created to handle motion, but cannot handle eye blink motion – as this does require some sort of indirect perception and representation to store old visual information, and possibly infer the location of moving information. We know of no models of filling-in across blinks or saccades based on optic flows. Current optic flow research investigates how the flow can be plausibly extracted (e.g., Arredondo, Lebart, & Lane, 2004; Chessa, Sabatini, & Solari, 2009), and how it can be put to ecological use such as disentangling object motion from self-motion (P. A. Warren & Rushton, 2009), segregating and tracking objects

(Farrajota, Rodrigues, & Du Buf, 2011), and for heading or steering control (Browning, Grossberg, & Mingolla, 2009; W. K. Page & Duffy, 2008).

C.2 Emergic model

In order to explain the phenomena of this chapter, we have taken the Emergic Cognitive Model (ECM) architecture as summarized on page 5, and made no changes. For the filling-in across eye blinks, we have one intuition as to how ECM might explain the results.

During an eye blink, information from the photoreceptors is supressed and not delivered to the first level RFs. The *handling missing data* function simply uses lateral information if bottom-up information is not available. If the eye is moving (or not) the *maintaining information coherence* function performs the lateral shifting of information. It is lateral recurrence that provides an emergic "memory". It is not what one would ordinarily conceive as actual memory, as it is not associated with any retinotopic neuron and it changes from tick to tick. It is actually associated with the spatiotopic continuously changing *flow* of information.

C.2.1 Blink Generation

Originally, blinks were simulated by detaching a *person* from the *world view* and re-attaching them afterwards. This effectively prevents visual information from reaching the photoreceptors. We have since added two behavioural parameters to automatically control a *person's* blinking to the same effect. Such perfect periodicity is extremely unnatural, but suffices for our visual purposes. For natural blinking, further information is available for the spontaneous blink generator (Kaminer, Powers, Horn, Hui, & Evinger,

2011) and the relation to task and saccades (Fogarty & Stern, 1989; Goossens & Van Opstal, 2010).

C.3 Test 1: Simple surfaces (Love)

In this test, we demonstrate and characterize the emergic filling-in behaviour of the ECM architecture that occurs across eye blinks when information from the photoreceptors is suppressed. In particular, we measure the performance of filling-in over simple homogeneous chromatic surfaces – a set of block letters spelling "LOVE" – by comparison to the identical no-blink condition of Leibovitz (p60).

C.3.1 Methods

The methods used are identical to those of page 80 except as indicated herein.

C.3.1.1 Subject analysis summary

Same as in page 81 (Illustration 234), but with a repeated blinking behaviour every 8 ticks – eyes open for 5 ticks followed by eyes closed for 3 ticks. Originally, this
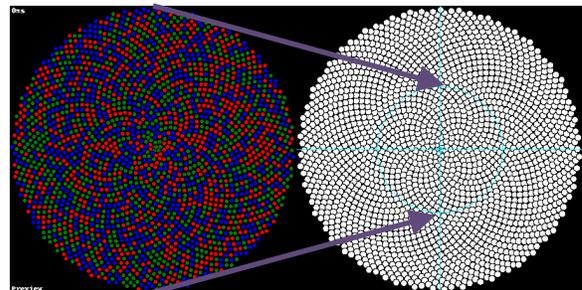
Illustration 234: Test 1 & 2 subject – their photoreceptor and RF1 mosaics

blinking behaviour was simulated within ESS by manually removing the *person* from the *world view*, and then reconnecting them. However, this was simplified by a set of blink related behavioural parameters and thus becomes a detail of *person* rather than procedure. These details are available in the supplementary material (s377).

C.3.1.2 Stimuli

The same stimulus as on page 81 was used and is shown in Illustration 235 (see also s379).

Illustration 235: Test 1 stimulus with view overlays

C.3.1.3 Procedure

The same procedures as on page 82 was used, but with the addition of eye blink. This was originally simulated within ESS by detaching the *person* from the *world view* and then reattaching them. However, this has been simplified by the addition of *person* behavioural parameters. The *person* was "instructed" to blink repeatedly in the following pattern

- eyes open for 5 ticks, then
- eyes closed for 3 ticks

while saccading over the stimuli in the following pattern.

- Move right for 8 ticks, target=(60′, 0′)
- Move down for 8 ticks, target=(60′, -60′)
- Move left for 8 ticks, target=(0′, -60′)
- Move up for 8 ticks, target=(0′, 0′)

The artificial brain was monitored for the first 32 time ticks. Note that there is an eye movement planning and distribution delay before the eye actually moves as desired.

C.3.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s382). When presented as an animation (Leibovitz, 2012a), all the results for this section can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=308
```

In this section, we summarize and discuss various highlights. However, we do not extract quantitative results, but qualitatively compare the filling-in behaviour under the blink condition, with the same behaviour under the no-blink condition of Leibovitz (p60).

C.3.2.1 Filling-in behaviour

The dynamic behaviour of filling-in during blink suppression is summarized in Illustration 236. Labels in red represent blink intervals for the photoreceptors. Therefore the level 1 receptive fields (RFs) are impacted one tick afterwards. The central region, bounded by the cyan circle, contains those RFs with bottom-up connections from photoreceptors (but do not receive information during an eye blink). Beyond this retinal region, the RFs can only receive information from lateral RFs, if the level has shifted to compensate for eye motion. Shifting is indicated via grey label backgrounds and starts at tick 7. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.

t=1	t=2	t=3	t=4	t=5	t=6	t=7	t=8
t=9	t=10	t=11	t=12	t=13	t=14	t=15	t=16
t=17	t=18	t=19	t=20	t=21	t=22	t=23	t=24
t=25	t=26	t=27	t=28	t=29	t=30	t=31	t=32
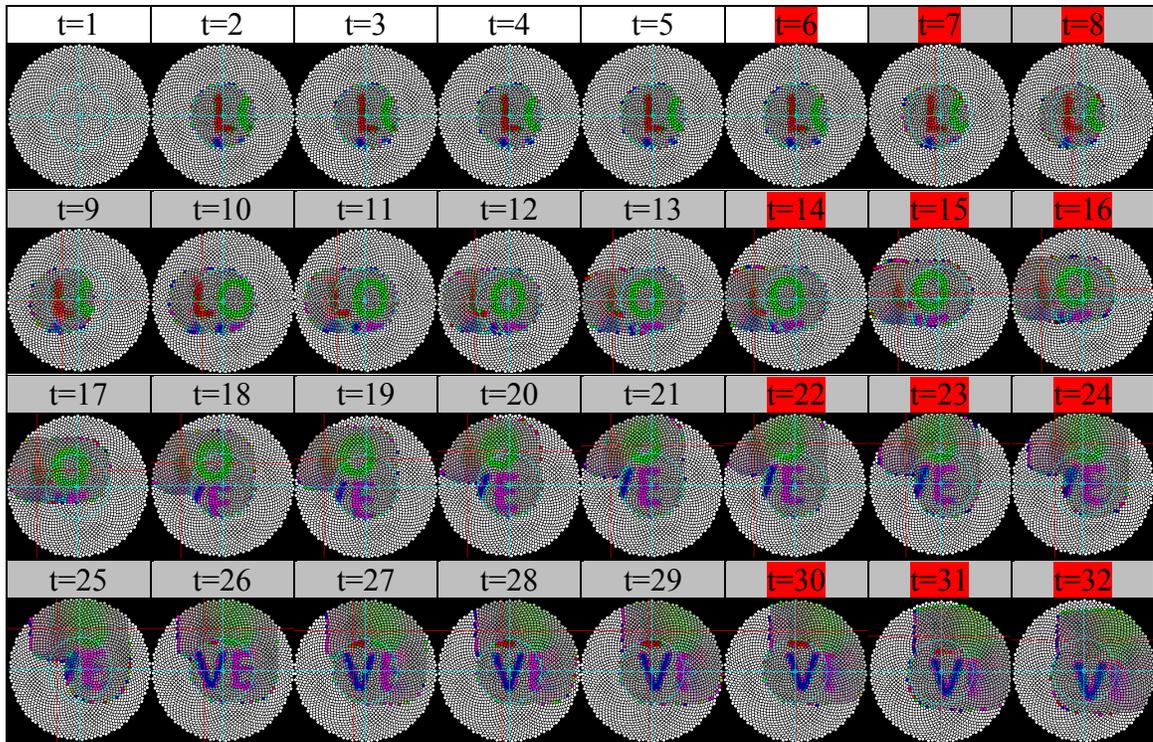
Illustration 236: Filling-in RF level 1 with simple surfaces shown in retinocentric coordinates

The behaviour for the level 2 and level 3 RFs is similar and therefore not shown here. See the supplementary material for details.

C.3.2.2 Filling-in comparison between blink and no-blink condition

While filling-out *beyond* the retina is demonstrated whenever the eyes move, the filling-in *within* the retina is demonstrated during eye blinks. In Illustration 237 we compare the blink condition of this test with the corresponding no-blink condition of Leibovitz (p60).

t=23	t=24	t=25	t=26

Blink Condition
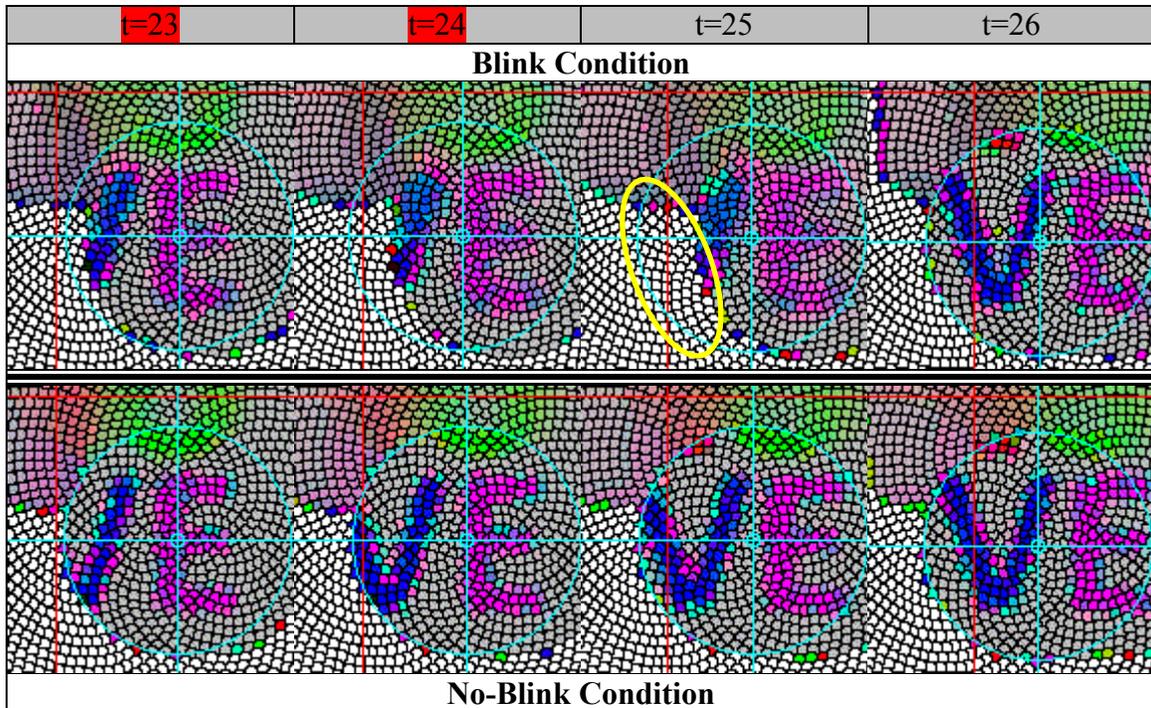
No-Blink Condition

Illustration 237: Filling-in of RF level 1 in blink and no-blink conditions

C.3.3   Discussion

It is difficult to detect the presence of a blink at the RF level 1 via cursory inspection of the filling-in behaviour (Illustration 236). This is exactly what is required by filling-in. However, a detailed examination of tick 25 does reveal minor differences (Illustration 237). While the letter 'E' is still cohesive, it has started to dissipate. At this particular time, the eye is moving leftwards, so the information is being shifted rightwards to compensate. More than half of the South-West quadrant has been filled-in with no-data (shown in white) so that the letter 'V' does not get completed. When the eye reopens at tick 26, the 'V' is properly re-formed and the blink condition becomes equal to the no-blink condition. The North-West quadrant isn't as problematic as there happens to be some information that can be shifted in.

The blink and no-blink conditions are identical with respect to the region beyond the retina which is always populated by information shifted in, so it is immune to eye blinks. During an eye-blink, the retinal and extra-retinal regions exhibit the same behaviour and underlying mechanism, so the terms *filling-in* or *filling-out* are merely descriptive. During the eye-blink, it makes no sense to even talk about the retinotopic reference frame for neurons, as they would have to be black. The flow of information is always spatiotopic, and the RFs merely sample the flow retinotopically, whether the flow comes up from the photoreceptors, laterally from others RFs, or from above. The *maintain information coherence* function operates continuously as does the *handle missing information* function. However, the latter behaviourally bifurcates during an eye blink giving lateral flows precedence over bottom-up flows.

C.4   Test 2: Natural Image (Lena)

In this test, we demonstrate and characterize the emergic filling-in behaviour of the ECM architecture that occurs across eye blinks when information from the photoreceptors is suppressed. In particular, we measure the performance of filling-in over a standard natural image – that of Lena (Picard, 1995) – by comparison to the identical no-blink condition of Leibovitz (p60).

C.4.1   Methods

The methods used are identical to Test 1 except as detailed herein.

C.4.1.1 Stimuli

The same stimulus as on page 81 was used here and is shown in Illustration 238 (see also s379).

C.4.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary
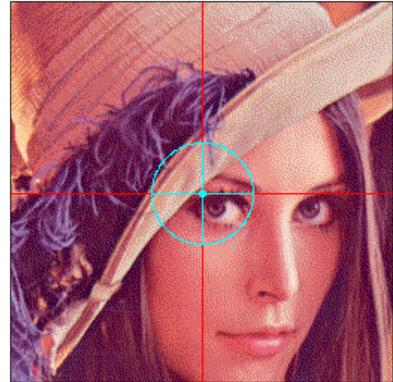
Illustration 238: Test 2 stimulus with view overlays

material (s393). In this section, we summarize and discuss various highlights. In this chapter we do not extract quantitative results, but qualitatively compare the filling-in behaviour under the blink condition, with the same behaviour under the no-blink condition of Leibovitz (p60).

C.4.2.1 Filling-in behaviour

The dynamic behaviour of filling-in during blink suppression is summarized in Illustration 239. Labels in red represent blink intervals for the photoreceptors. Therefore the level 1 RFs are impacted one tick afterwards.
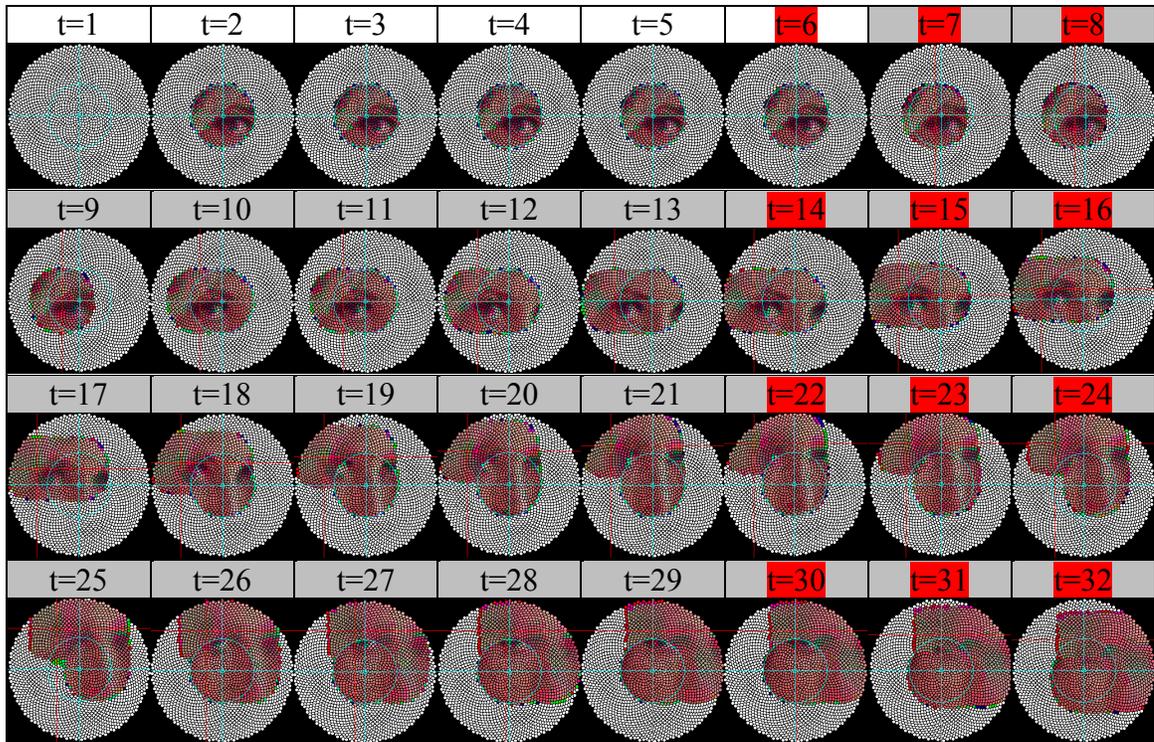
Illustration 239: Filling-in RF level 1 with natural image shown in retinocentric coordinates

The behaviour for the level 2 and level 3 RFs is similar and therefore not shown here. See the supplementary material for details.

C.4.2.2 Filling-in comparison between blink and no-blink condition

While filling-out *beyond* the retina is demonstrated whenever the eyes move, the filling-in *within* the retina is demonstrated during eye blinks. In Illustration 240 we compare the blink condition of this test with the corresponding no-blink condition of Leibovitz (p60).

t=7	t=8	t=9	t=10

Blink Condition
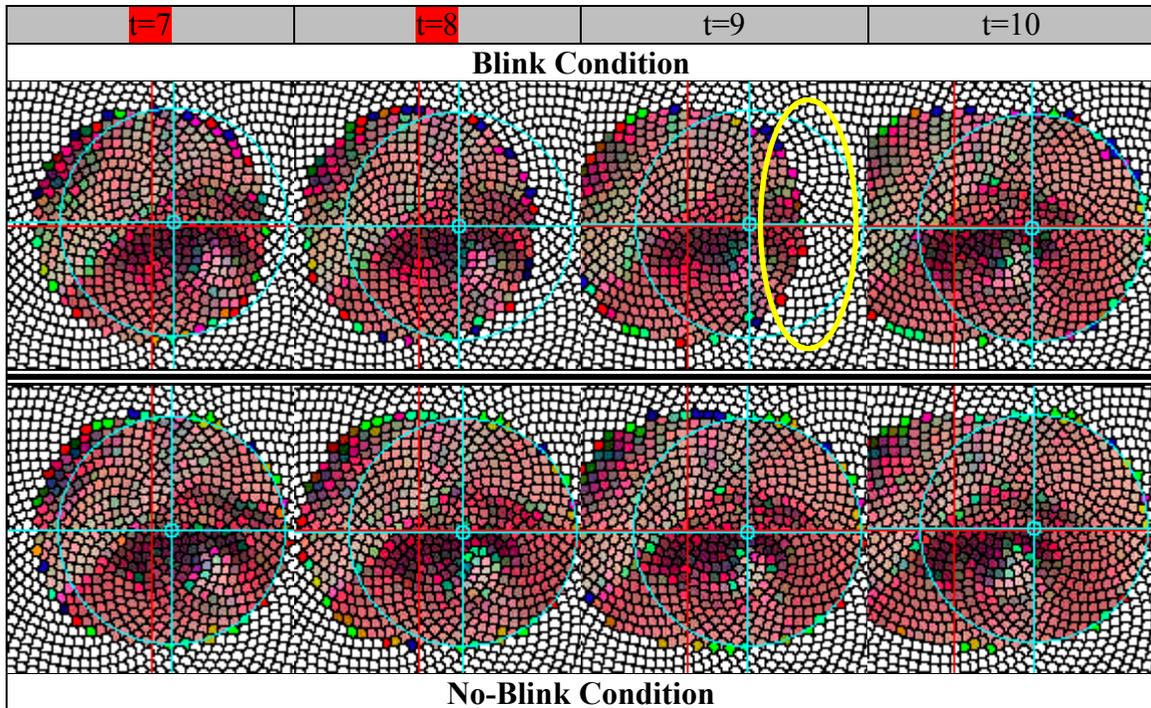
No-Blink Condition

Illustration 240: Filling-in of RF level 1 in blink and no-blink conditions

C.4.3 Discussion

It is difficult to detect the presence of a blink at the RFs via cursory inspection of the filling-in behaviour (Illustration 239). This is exactly what is required by filling-in. However, a detailed examination of tick 9 does reveal minor differences (Illustration 240). Lena's eye remains cohesive during the entire blink suppression period and it has lost very little detail as compared to the letter 'E' of Test 1. At this particular time, the agent's eye is moving rightwards, so the information is being shifted leftwards to compensate. More than half of the North-East and South-East quadrants have been filled-in with no-data (shown in white) so that Lena's eye is not extended (filled-out) as required. When the agent's eye reopens at tick 10, then Lena's eye is properly extended and the blink condition becomes equal to the no-blink condition.

C.4.3.1 Comparison to alternative model

We repeat an afterimage model (Momiji et al., 2006) as shown in Illustration 241 that was mentioned for flash memories (p13). While these dynamics may have been acceptable when the image was permanently removed, they are not acceptable during a blink period – images
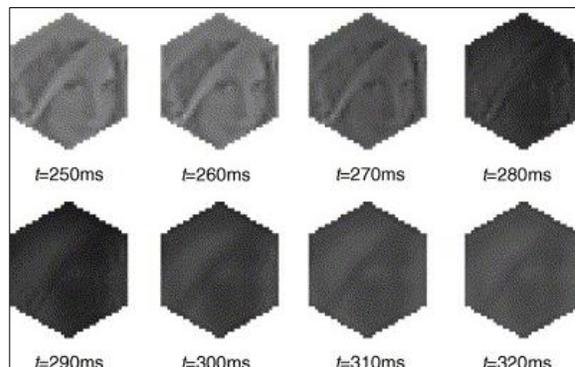
Illustration 241: Afterimages as memories.
From (Momiji et al., 2006), Fig. 16. Stimulus fixated for 250ms then blanked.

do not appear to fade and reverse during a blink which is not salient at all. This is not an actual problem with the original model as it was not intended to handle blinks. It was also not intended to handle eye motion during blink, so if the eye re-opens in this model, there would be a jarring disconnect between the unshifted afterimage and the shifted visual stimulus. However, ECM handles both flash memories and blinks with aplomb and without a change in parameters – indicative of a more useful functional decomposition.

C.5 Test 3: Dynamic image with surface gradients (Chaser)

In this test, we demonstrate and characterize the emergic filling-in behaviour of the ECM architecture that occurs across eye blinks when information from the photoreceptors is suppressed. In particular, we measure the performance of filling-in over a dynamic image that contains surface gradients – a Lilac Chaser – by comparison to the identical no-blink condition of Leibovitz (p60).

C.5.1 Methods

The methods used are identical to Test 1 except as detailed herein.

C.5.1.1 Subject analysis summary

The same agent as in Test 1 was used except that the agent was positioned further away from the virtual computer monitor and this behavioural parameter happens to be specified within *person*. The detailed parameters are available in the supplementary material (s377).

C.5.1.2 Stimuli

The same stimulus as on page 81 was used and is shown in Illustration 242 (see also s380).

C.5.1.3 Procedure

The same procedure as in Test 1 was used except that movement was in 40′ increments rather than 60′.
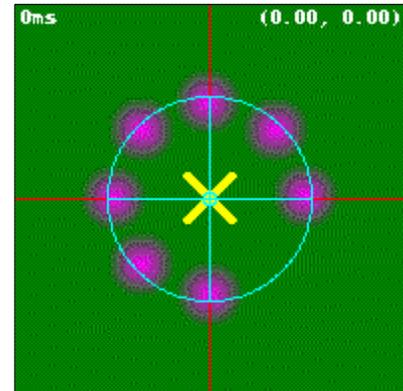
Illustration 242: Test 3 stimulus with view overlays

C.5.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s404). In this section, we summarize and discuss various highlights. In this chapter we do not extract quantitative results, but qualitatively compare the filling-in behaviour under the blink condition, with the same behaviour under the no-blink condition of Leibovitz (p60).

C.5.2.1 Filling-in behaviour

The dynamic behaviour of filling-in during blink suppression is summarized in Illustration 243. Labels in red represent blink intervals for the photoreceptors. Therefore the level 1 RFs are impacted one tick afterwards.
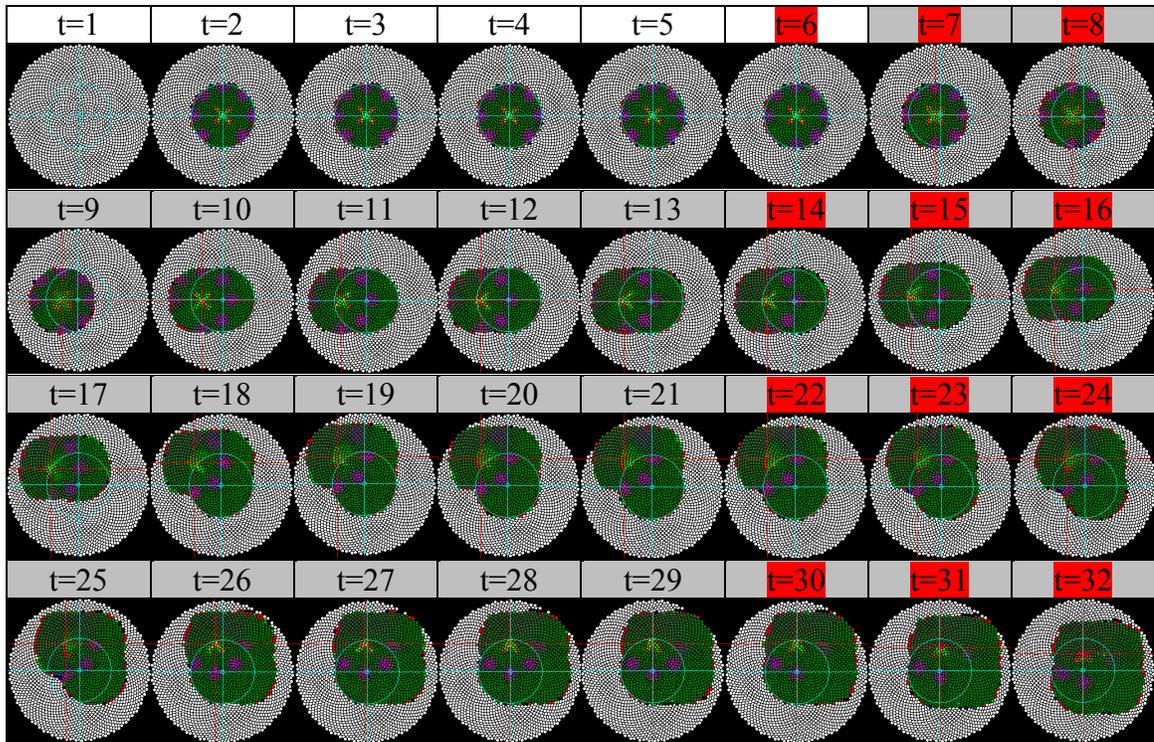
Illustration 243: Filling-in RF level 1 with dynamic image shown in retinocentric coordinates

The behaviour for the level 2 and level 3 RFs is similar and therefore not shown here. See the supplementary material for details.

C.5.2.2 Filling-in comparison between blink and no-blink condition

While filling-out *beyond* the retina is demonstrated whenever the eyes move, the filling-in *within* the retina is demonstrated during eye blinks. In Illustration 244 we compare the blink condition of this test with the corresponding no-blink condition of Leibovitz (p60).

t=15	t=16	t=17	t=18

Blink Condition
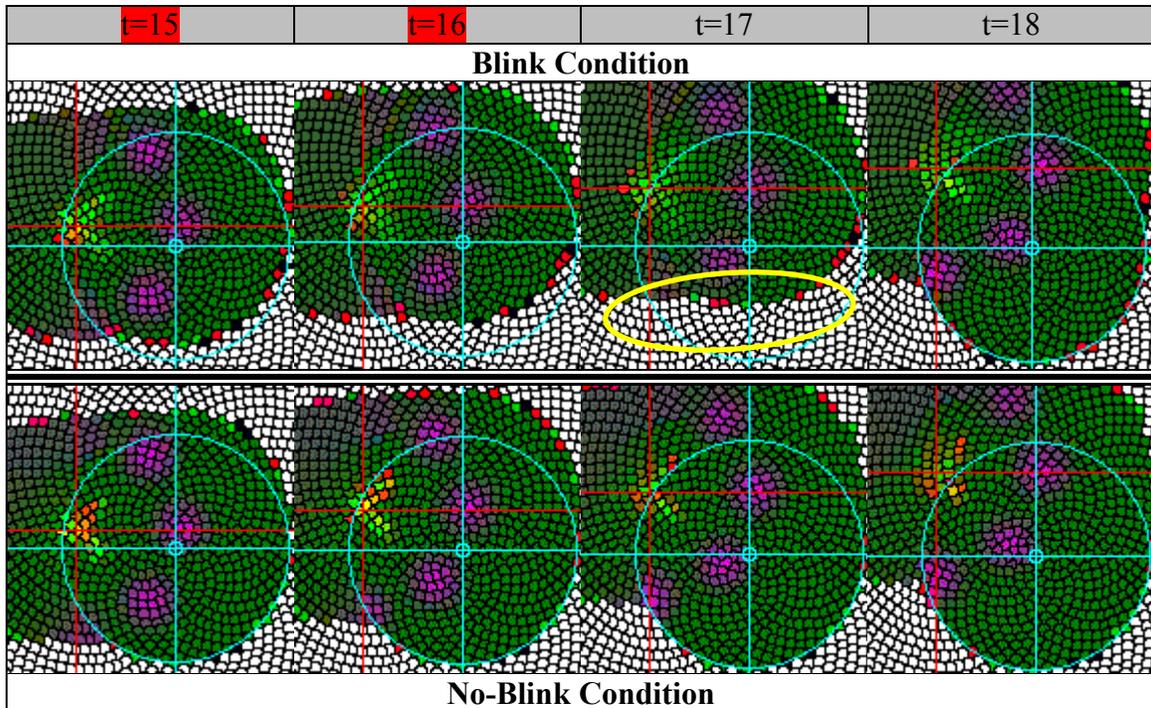
No-Blink Condition

Illustration 244: Filling-in of RF level 1 in blink and no-blink conditions

C.5.3 Discussion

It is difficult to detect the presence of a blink at the RF level 1 via cursory inspection of

the filling-in behaviour (Illustration 243). This is exactly what is required by filling-in.

However, a detailed examination of tick 17 does reveal minor differences (Illustration

244). The lilacs remain cohesive during the entire blink suppression period and have just

lost a little bit of their central brightness. They have lost very little detail as compared to

the letter 'E' of Test 1. At this particular time, the agent's eye is moving downwards, so

the information is being shifted upwards to compensate. More than half of the South-

West and South-East quadrants have been filled-in with no-data (shown in white) so that

the new lilac coming in is only $1/3^{rd}$ formed. When the agent's eye reopens at tick 18,

then all the lilacs are properly refreshed, and the blink condition becomes equal to the no-

blink condition.

C.6 Test 4: Rapid and extended blinking (Square)

In this test, we demonstrate and characterize the emergic filling-in behaviour of the ECM architecture that occurs across rapid and extended eye blinks when information from the photoreceptors is suppressed. In particular, we measure the performance of rapid filling-in over a homogeneous chromatic surface – a square.

C.6.1 Methods

The methods used are identical to Test 1 except as detailed herein.

C.6.1.1 Subject analysis summary

Same as Leibovitz (p103) and shown in Illustration 245, but with a repeated blinking behaviour every 3 ticks – eyes open for 1 tick followed by eyes closed for 2 ticks. Originally, this blinking behaviour was simulated within ESS by manually detaching the *person* from the *world*

Illustration 245: Test 4 subject – their photoreceptor mosaic

view, and then reattaching them. However, this was simplified by a set of blink related behavioural parameters and thus becomes a detail of *person* rather than procedure. These details are available in the supplementary material (s378).

C.6.1.2 Stimuli

The same stimulus as on page 220 was used and is shown in in Illustration 246 (see also s380).

C.6.1.3 Procedure

A similar procedure as Test 1 was used but with faster eye

Illustration 246: Test 4 stimulus with view overlays

blinks and faster back and forth motion. The *person* was "instructed" to blink repeatedly

in the following pattern

- eyes open for 1 tick, then

- eyes closed for 2 ticks

while saccading over the stimuli in the following pattern.

- Move right at 2.5′/tick for 16 ticks

- Move left at 2.5′/left 22 tick

The artificial brain was monitored for the last 32 time ticks. Note that there is an eye

movement planning and distribution delay before the eye actually moves as desired.

C.6.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed

in the supplementary material (s415). In this section, we summarize and discuss various

highlights.

C.6.2.1   Filling-in behaviour

The dynamic behaviour of filling-in during blink suppression is summarized in

Illustration 247. Labels in red represent blink intervals for the photoreceptors. Therefore

the level 1 RFs are impacted one tick afterwards.
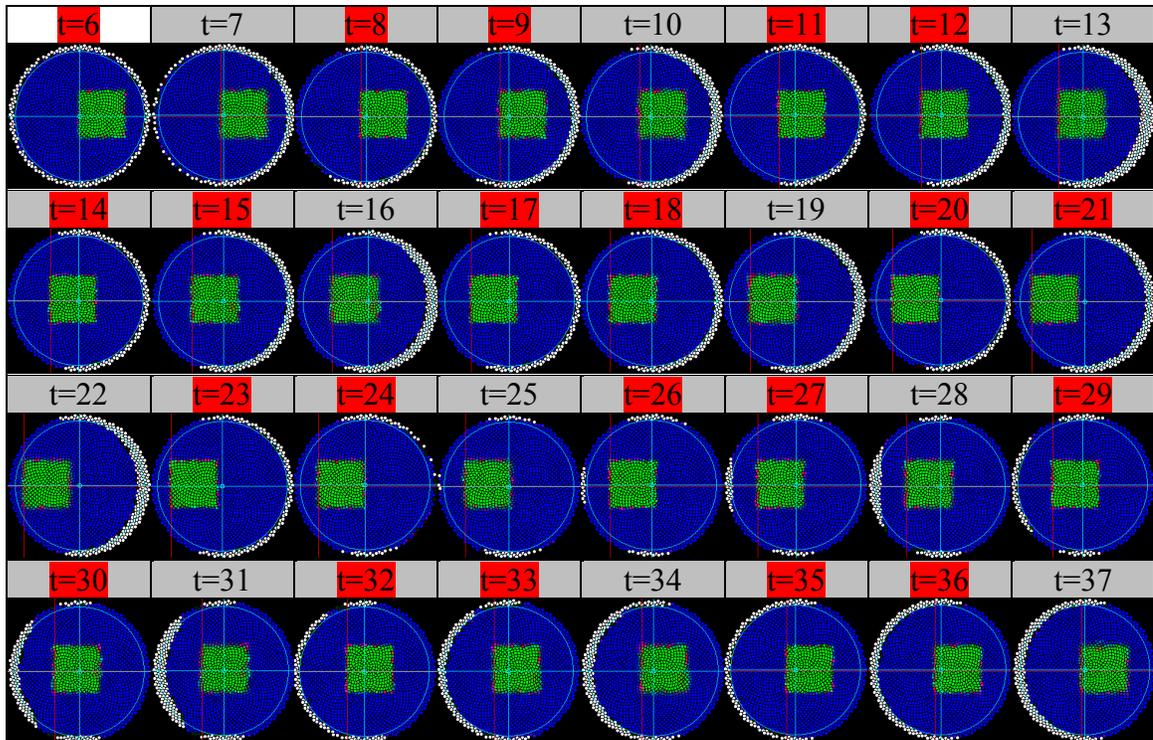
Illustration 247: Filling-in RF level 1 of square image shown in retinocentric coordinates

The behaviour for the level 2 and level 3 RFs is similar and therefore not shown here. See the supplementary material for details.

C.6.2.2 Filling-in comparison between blink and no-blink condition

While filling-out *beyond* the retina is demonstrated whenever the eyes move, the filling-in *within* the retina is demonstrated during eye blinks. In Illustration 248 we compare the temporal aspects of filling-in across a single blink period with the no-blink condition of (Leibovitz, 2012a) in greater detail.

t=9	t=10	t=11
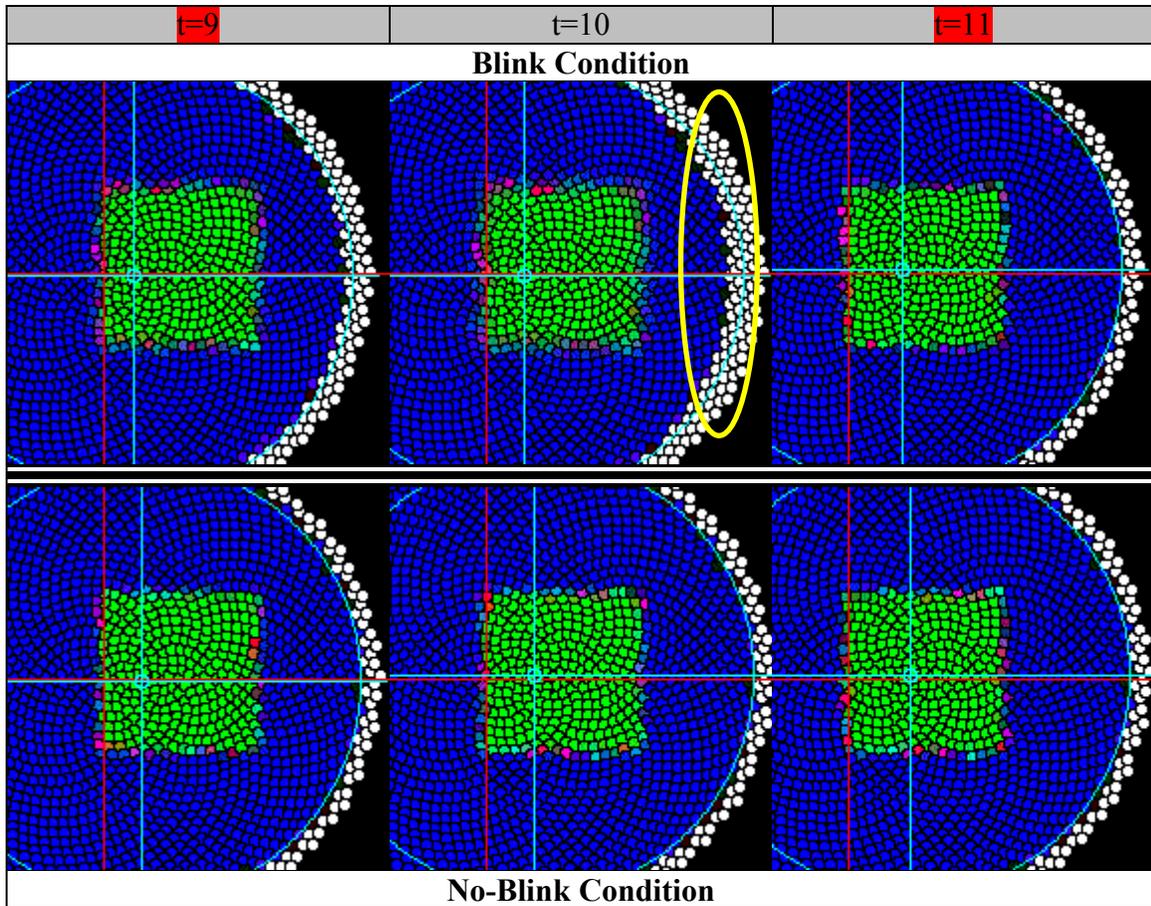
Illustration 248: Filling-in of RF level 1 in blink and no-blink conditions

C.6.2.3 Smooth Motion

With the eyes open only 1/3rd of the time, the photoreceptors sense a jumpy motion from one open blink to another as shown in Illustration 249. Nevertheless, as shown in Illustration 248, but much more salient in the animations (Leibovitz, 2012a) available at http://emergic.upwize.com/?page_id=308, motion appears smooth while the eyes are closed.

t=13	t=16	t=19
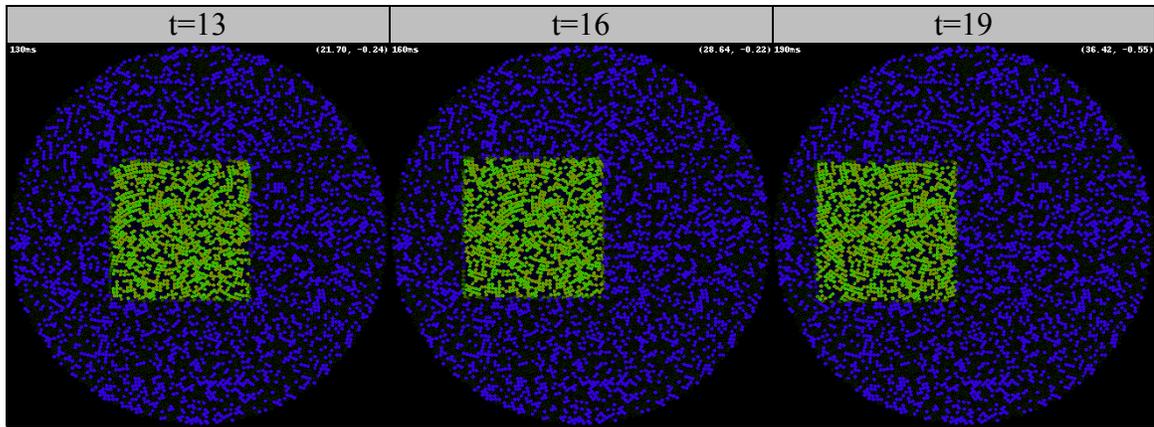
Illustration 249: Jerky motion across open eye periods

C.6.3   Discussion

It is difficult to detect the presence of a blink at the RF level 1 via cursory inspection of the filling-in behaviour (Illustration 247). This is exactly what is required by filling-in. However, a detailed examination of tick 10 does reveal minor differences (Illustration 248). The borders of the green square remains cohesive during the entire blink suppression period. They have lost very little detail as compared to the letter 'E' of Test 1. At this particular time, the agent's eye is moving rightwards, so the information is being shifted leftwards to compensate. A minor portion of the North-East and South-East quadrants have been filled-in with no-data (shown in white), but when the agent's eye reopens at tick 11, then all the missing retinal information is restored.

Motion appears smooth as it is internally generated in minimal amounts per tick, and the "internal model" is effectively updated every tick, regardless of whether eyes are open or closed. However, this phenomenon is not the same as the smooth motion perceived when objects move, possibly across occlusions and including the inter-stimulus intervals

between movie frames. We have not yet implemented a motion detection system, but this phenomenon demonstrates that the basic infrastructure is available.

C.7   General results and discussion

C.7.1   Results summary

Visual information is filled-in when the eyes are closed and information from the photoreceptors is supressed. If the eye happens to be saccading at this time, then the "internal model" is smoothly updated. The fidelity of the image is less affected by the percentage of time the eye is closed, and more by the amount of time it is closed per blink period. For example, in the rapid and extended eye blink test, the eye is closed $2/3^{rd}$ of the time, but the closed duration is only 2 ticks. In the other tests, the eye is closed $3/8^{th}$ of the time, but the closed duration is 3 ticks – one tick longer.

Natural images degrade less during blinks. A single homogeneous square does not degrade as much as more complicated yet still homogeneous coloured surfaces such as letters.

C.7.2   Comparison to alternative model

Billock (1997) proposed that the recurrence between the LGN and V1 cortex can set up a *reverberation* that acts as a very short term visual memory supporting the temporal filling-in across blinks (and saccades). Such reverberation would have to be inhibited during ordinary vision. No computational model was associated with this conceptual "theory". ECM has no reverberations between levels that could function as memory. It does have recurrent lateral connectivity that could in principle cause reverberations. During an eye blink, and assuming the eye does not move, each RF effectively sends

information back to itself. While the eye is open, this also occurs in the region beyond the photoreceptors. However, no recurrence or reverberations are required. As long as the eye is moving in one direction, then the information can flow laterally in the opposite direction – connectivity could be in one direction without looping back. Indeed, as long as the eye is moving in a non-reverberating pattern, the information flows laterally in an equal and opposite non-reverberating pattern. Moreover, due to neural heterogeneity, the information from two incoming RFs can be combined, or the information from a single RF can be split, as the flow of information traverse the RFs laterally. In any case, there is no maintenance of the same data as would be stored by one neuron, but there would be an effective memory within the flow.

C.7.3 Sensation vs. perception

We are extremely careful in distinguishing sensation from perception. To us, perception is mostly cognitively impenetrable (Pylyshyn, 1999b) and we do not model any of its aspects such as colour qualia. Nevertheless, all perceptual phenomena do shed light on the underlying sensory apparatus which we do model in a functional manner. Because it is never clear cut as to how much of any one phenomenon can be attributed to perceptual or sensory mechanisms, our advice is never to model phenomenon explicitly as more often than not, they will emerge naturally due to the interaction of finer grained sensory systems.

We view the sensory world is an autonomous, massively parallel system designed to provide value added service to the attentionally limited perceptual world. Why would the sensory system keep a detailed model of the world up to date even across eye blinks? We

believe that this is to enable the sensory system to supply arousal (attend-to-this) signals to perception. These include all sensory changes that are surprising and differ from predictions.

For example, while all objects disappear during an eye blink, they are maintained including any change in position, so that when the eye re-opens, the sensory system can notify perception of interesting changes that occurred because the world changed, and not because the eye blinked. Mitroff & Scholl (2004) found cases where objects enter conscious awareness for the first time only after they have disappeared. Not only does this imply the existence of a surprise signal, but it also implies a non-representational visual afterimage of the deleted object – "you can see the disappearance of something you can't see"!

However, prediction during an eye blink is not expected to be perfect, so some extra leeway is given before surprise in minor object position changes is declared. But, if a target of a saccade is blanked and restored 50-300ms afterwards, then even minor displacement changes can be detected – displacement detection during exogenous blanking differs from endogenous blinking (Deubel et al., 2004). On a similar vein, Domhoefer et al. (2002) found that change detection during blinks and blanks outperforms saccades but only in dynamic scenes. This implies that models must take into account motion, as does ours.

C.8 Conclusion

In the ECM architecture, whenever information is suppressed from one flow, i.e., bottom-up from the photoreceptors, then it is utilized from another, e.g., lateral flow. Thus, the suppression of visual information during eye-blinks is handled automatically in ECM. Moreover, the flows in ECM are always spatiotopic and take into account eye motion. Thus, during an eye closed interval, the visual information is effectively interpolated smoothly across time. If the eye closed duration is not overly long, the degradation in image quality is not noticeable, especially for natural images. The filling-in to support blink suppression is a robust mechanism effective against a range of stimuli. We hypothesize that this sensory infrastructure underlies an attentionally limited perceptual system.

Appendix D: An emergic model of filling-in from imagination

Humans dream at night and imagine during the day. Some imaginations are to perform specific tasks such as mental rotation. Some are to construct non-visual targets such as the center of an object, or to predict where an opening will be as two lines of marchers in a band cross each other. There are at least three cognitive
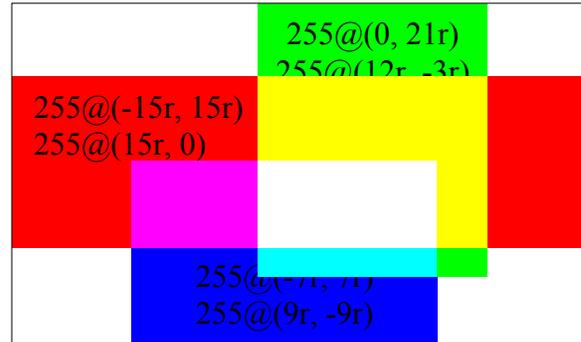
Illustration 250: Imagined rectangles.
Three Emergic Values delivered once to the three RGB (LMS) *sensory* pathways of the central RF at Level 3. Intersecting colours shown on GUI with no *perceptual* support within ECM. r=radius of central RF; coordinates represent statistics within Emergic Values (Top Left and Bottom Right corners).

reasons to continually imagine or generate a scene. 1) To understand a scene we may need to be able to create it (Hinton, 2007). This has an evolutionary advantage – rather than evolve two separate systems in parallel, an image decomposition process and an image composition process with some common representational interchange format, one only need evolve a single image generation mechanism – if you can generate the original image, then you have fully understood it (see feedback decoding in Swindale, 2008). 2) The only purpose of learning/induction is to use past experience to improve future performance (Kveraga, Ghuman, & Bar, 2007) which leads to an enactive view of memory (Klein, Robertson, & Delton, 2010). Thus, the only reason we may have episodic memory (any memory) is to enable future thinking. Both episodic recall and episodic future thinking may be subserved by the same neurocognitive process involving sensory and spatial imagery (Berntsen & Bohn, 2010; Byrne, Becker, & Burgess, 2007). 3) Visual percepts may be predictions of the future to enable us to catch things with our delay filled brain (Cudeiro, 2008; Cunningham, 2008; Muckli, Vetter, & Smith, 2011;

Nijhawan & Wu, 2009; Perrett et al., 2009; Roach, McGraw, & Johnston, 2011). Perhaps this helps to explain image stability and the predictive remapping of RFs?

Such views can turn bottom-up thinking on its head, and reverse the driver/modulation distinction (Sherman & Guillery, 1998). The important flow of information is top-down prediction and control which is modulated by bottom-up confirmation or disapproval in the form of surprises – deviations from predictions. (Neisser's (1976) Perceptual Loop would be an example of such an architecture. Indeed, once motion is taken into account, the *significant* change that occurs during the remapping of receptive fields cannot be thought of as *modulatory*.

The Emergic Cognitive Model has a single flow which can be (sometimes arbitrarily) subcategorized into bottom-up, top-down and lateral. None of these has superiority over another. They all influence each other, and emergic computations are nothing more than formalizing the influence relations. The example system of equations $a \mathrel{+}= f_a(a,b,c)$ and $b \mathrel{+}= f_b(a,b,c)$ demonstrate this. Neither a nor b is superior. Nevertheless, it is sometimes worthwhile to study the system in one direction to increase understanding in some epistemic fashion. In visual perception, bottom-up (or feedforward) analysis has predominanated, but in the rest of cognitive science, computationalism has been a top-down affair. ECM unifies the two.

All our computational models have so far been initially driven by visual stimuli. Certainly memory effects are due to lateral connectivity, and filling-in effects can be due

to top-down recurrence. However it is sometimes worthwhile to see the system in action within a different mode of operation. In this chapter, we present a cognitive model of filling-in from imagination. The purpose is to epistemically play with the model and characterize its behaviour. It provides another view into how structured information values are represented, and how this impacts the duration of memories-in-motion. In particular, it shows the power of recurrence and the possible economy in scene representation.

D.1 Test 1: Imagine quickly

A single sensory "thought" is injected top-down into the system, and in this chapter we characterize its behaviour.

D.1.1 Methods

The methods used are identical to those of page 36 except as indicated herein.

D.1.1.1 Subject analysis summary

The same *person* from page 81 was used. During imagination, it does not matter whether RFs correspond to retinal or extra-retinal regions. The participant's retinal mosaic is shown in Illustration 251, while details can be found in the supplement (s426). Note that imagination occurs within the RFs and not within the retina.
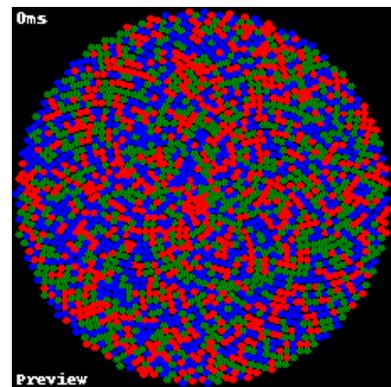
Illustration 251: Imagination Test 1 (Fast) agent

D.1.1.2 Stimuli

A single sensory "thought" (Illustration 250) is injected into the central RF of the top-most RF level. This "thought" consists of an L, M and S set of emergic values, of type

ValueXY representing statistical information with spatial extent. The emergic value

ValueXY has the following signature

```
ValueXY((x,y,v,w)…)
```

where

      x: the x-coordinate of this sample value

      y: the y-coordinate of this sample value

      v: the L, M or S colour value of this sample from 0 to 255

      w: the weight assigned to this sample value

      …:   repeatable, in which case this emergic value represents the

                    descriptive statistics for all the samples weighted accordingly

If r is the radius of the central RF, then these are the three injected values

```
L/Red:     ValueXY((-15*r, 15*r, 255, 9),
                   ( 15*r,    0, 255, 9))

M/Green:   ValueXY((    0, 21*r, 255, 9),
                   ( 12*r, -3*r, 255, 9))

S/Blue:    ValueXY(( -7*r,  7*r, 255, 9),
                   (  9*r, -9*r, 255, 9))
```

Each value represents two opposite corners of a coloured rectangle.

Detailed information is found in the supplement (s428).

D.1.1.3   Procedure

At time=0, the three L, M and S emergic values were injected into the central RF in the

top-most level of *the virtual agent*. The thought was not repeated at other times. The

agent was instructed not to move its eye (barring involuntary jitter).

The artificial brain was monitored for the first 24 time ticks.

D.1.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s431). In this section, we summarize and discuss various highlights.

The frame-by-frame account for all 4 experiments encompasses more than 36 pages because animations cannot be printed. It is illustrated within tables that function as multi-track sequence diagrams annotating the dynamic "orchestra of the mind". When presented as an animation (Leibovitz, 2012a), all the results for this test, along with Experiments 2-4 can be displayed on a single web page.

```
http://emergic.upwize.com/?page_id=263
```

The animations exhibit behaviour, e.g., filling-in, and can indicate movement that is hard to detect across image frames, while the frame-by-frame account allows precise quantitative comparison of processing over time. Here we highlight some of the data in a more useful form under print technology.

D.1.2.1 Imagination filling-in behaviour

The model's behaviour of filling-in for imagination is summarized in Illustration 252. Eye jitter starts at tick 7 and is indicated via a grey title background. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
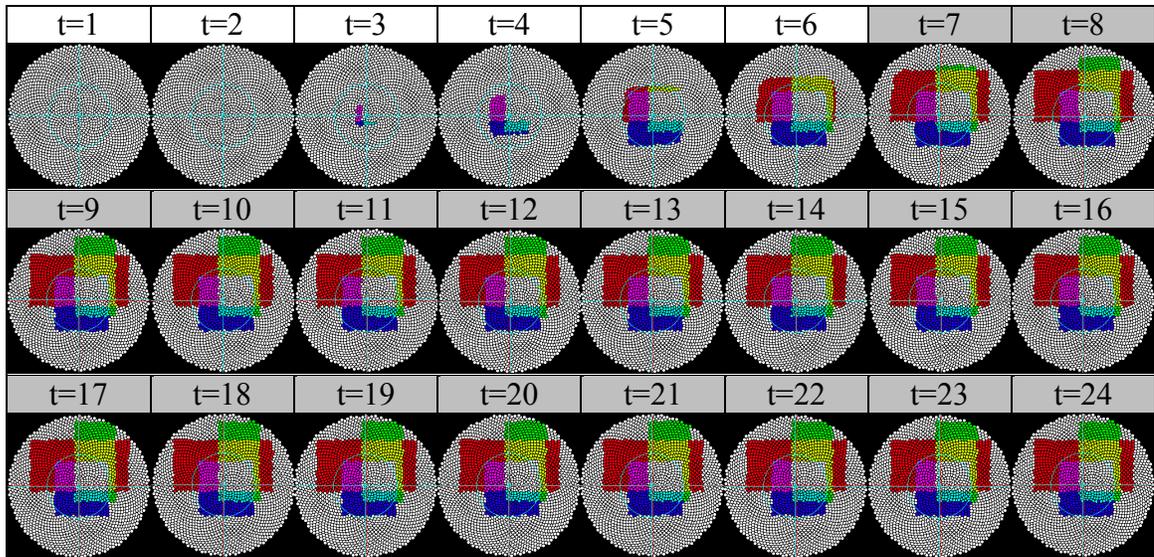
Illustration 252: Imagination filling-in behaviour of 1st Level RFs (Fast)

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These

details are available in the supplemental material.

Imagination filling-in completed by tick 9. From the perspective of the 1st level RFs, this

took 6 ticks to occur (from tick 3 to tick 9).

D.1.2.2 Imagination filling-in details

The detailed results for all three RF levels is shown in Illustration 253.

t=9 @ Level 1	t=9 @ Level 2	t=9 @ Level 3
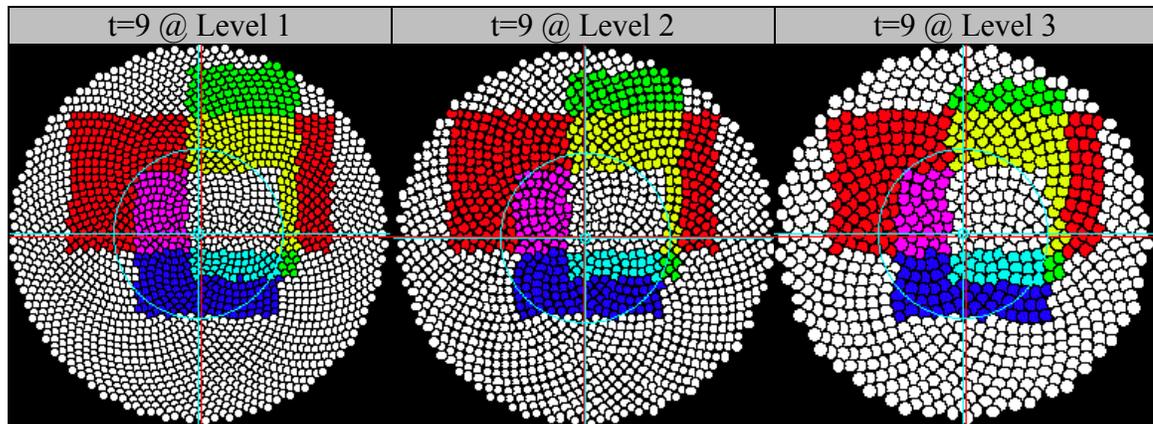
Illustration 253: Imagination filling-in details at tick 9 for all RF levels (Fast)

The borders are considered as straight within the tolerance of a single RF.

D.1.3 Discussion

We have demonstrated that a single "thought" or highly compact representation could cause a complex image to unfold over time.

D.1.3.1 Colour perception

The ECM model has separate L, M and S pathways throughout. At this point in time, they are not combined, partly because ECM is a *sensory* model and mostly because we do not yet have a story for conscious *perception*. Nevertheless, the GUI does combine L, M and S, but only to convert them to separate R, G and B colour coordinates which are still kept separate on the computer – a pixel is made up of 3 sub-pixels. While the graphical results looks quite complex having the following 7 colours: red, green, blue, yellow, cyan, magenta and white, the model only has three rectangles, one for each cone type. The complexity of the image is due to our human eyes, not to the model.

D.1.3.2 Straight lines

ECM has neither border contour nor edge processing; nevertheless the appearance of straight edges is an emergic property of the manner in which statistics are maintained with spatial extent.

D.1.3.3 Memory duration

The "thought" is injected once but the image does not decay over time. This would be an example of a memory that does not dissipate. This can be explained by the fact that there is only one emergic value that has no others to interfere with. As the value goes back up, it is averaged with copies of itself never changing the basic descriptive statistics.

D.1.3.4 Filling-in mechanism

The single injected Emergic Value is distributed down and back up, but most of the filling-in is due to lateral connectivity. The value is accepted by all RFs that happen to fall within the spatial extent specified in the value. In the next experiment, we will test whether changing lateral fan-out parameters will affect the filling-in duration.

D.2 Test 2: Imagine slowly

We wondered how reducing the lateral fan-out configuration would slow down the filling-in times, and in this test we characterize the behaviour of such a change.

D.2.1 Methods

The methods used are identical to Test 1 except as detailed herein.

D.2.1.1 Subject analysis summary

This agent was similar to that of Test 1 but with fan-out parameters reduced as detailed in the supplemental material (s427).

D.2.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s440). In this section, we summarize and discuss various highlights.

D.2.2.1 Imagination filling-in behaviour

The model's behaviour of filling-in for imagination is summarized in Illustration 254. Eye jitter starts at tick 7 and is indicated via a grey title background. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
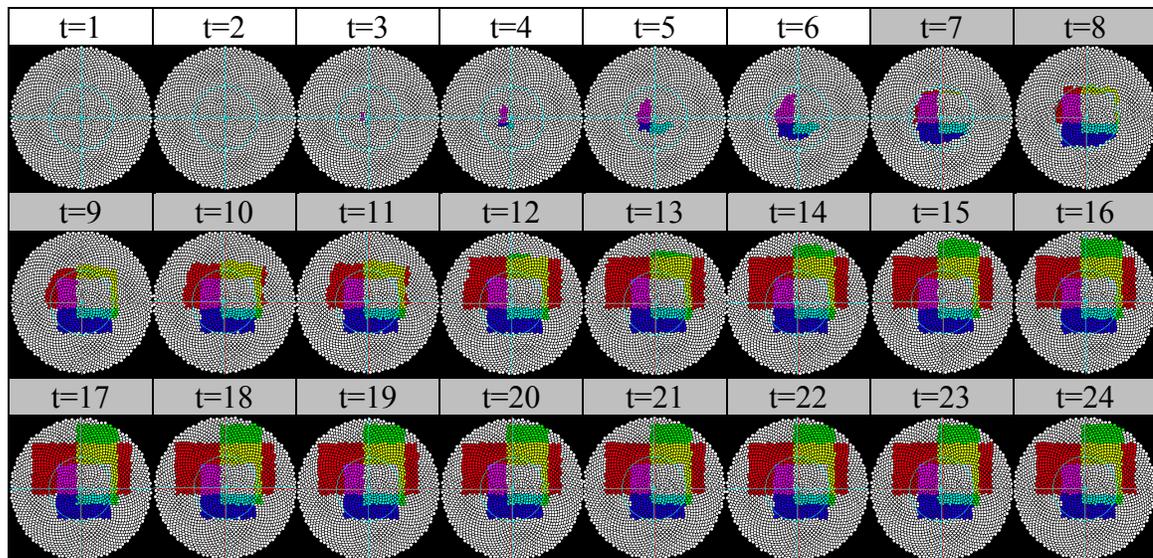
Illustration 254: Imagination filling-in behaviour of 1st Level RFs (Slow)

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

Imagination filling-in completed by tick 16. From the perspective of the 1st level RFs, this took 13 ticks to occur (from tick 3 to tick 16).

D.2.2.2 Imagination filling-in details

The detailed results for all three RF levels are shown in Illustration 255.
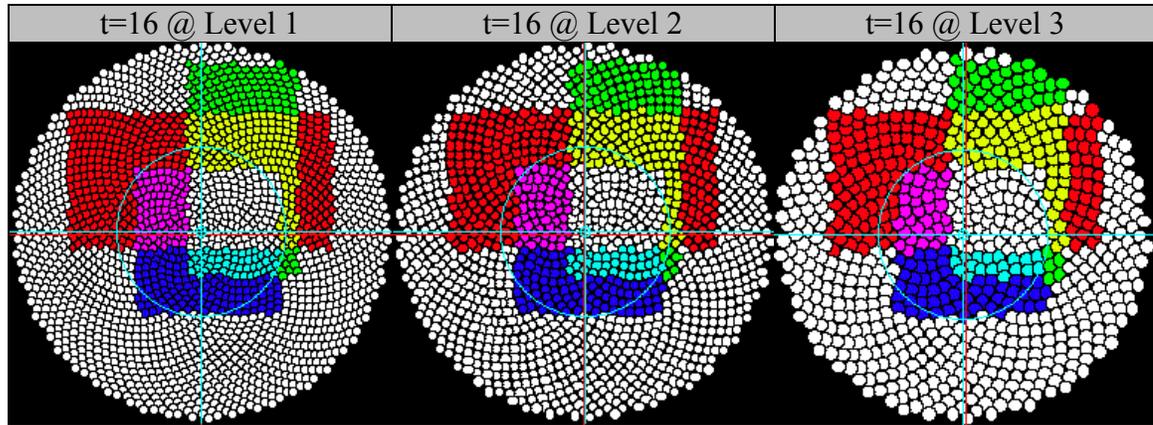
Illustration 255: Imagination filling-in details at tick 16 for all RF levels (Slow)

The borders are considered as straight within the tolerance of a single RF.

D.2.3 Discussion

With reduced lateral connectivity, filling-in times are twice as long. At tick 9, it is clear that information travels in a radiating pattern away from the central injection site of the "thought" – it does not travel along straight border lines.

D.3 Test 3: Imagine slowly with movement

In this test we wanted to characterize the filling-in of imagination behaviour during eye movements.

D.3.1 Methods

The methods used are identical to Test 2 except as detailed herein.

D.3.1.1 Procedure

The procedure was similar to Tests 1-2 but the *person* was instructed to saccade (while imagining) as follows

- Move right for 24 ticks at a speed of 2.5′/tick (± jitter)

The artificial brain was monitored for the first 24 time ticks. Note that there is an eye movement planning and distribution delay of 6 ticks before the eye actually moves as desired, so the above plan will not complete.

D.3.2   Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s449). In this section, we summarize and discuss various highlights.

D.3.2.1   Imagination filling-in behaviour

The model's behaviour of filling-in for imagination is summarized in Illustration 256. Eye jitter starts at tick 7 and is indicated via a grey title background. While every tick is nominally 10ms, we have not calibrated it with biology. At this point in time, we are only concerned with relative dynamics.
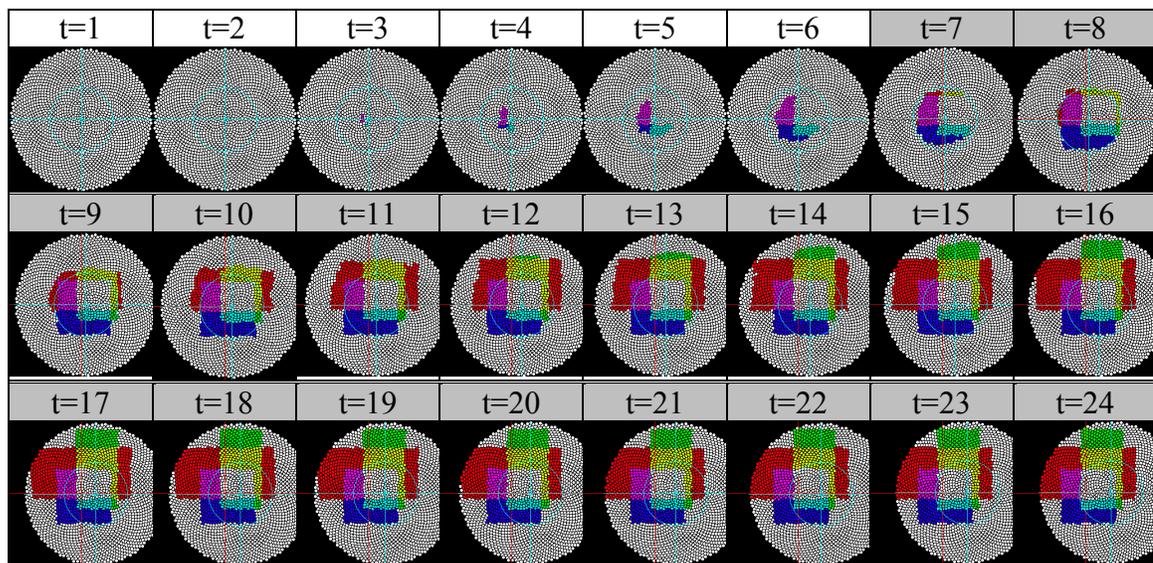
Illustration 256: Imagination filling-in behaviour of 1st Level RFs (Slow+Move)

Similar spatiotemporal filling-in behaviour occurs at the 2nd and 3rd RF levels. These details are available in the supplemental material.

Imagination filling-in completed by tick 16. From the perspective of the 1st level RFs, this took 13 ticks to occur (from tick 3 to tick 16).

D.3.2.2 Imagination filling-in details

The detailed results for all three RF levels are shown in Illustration 257.
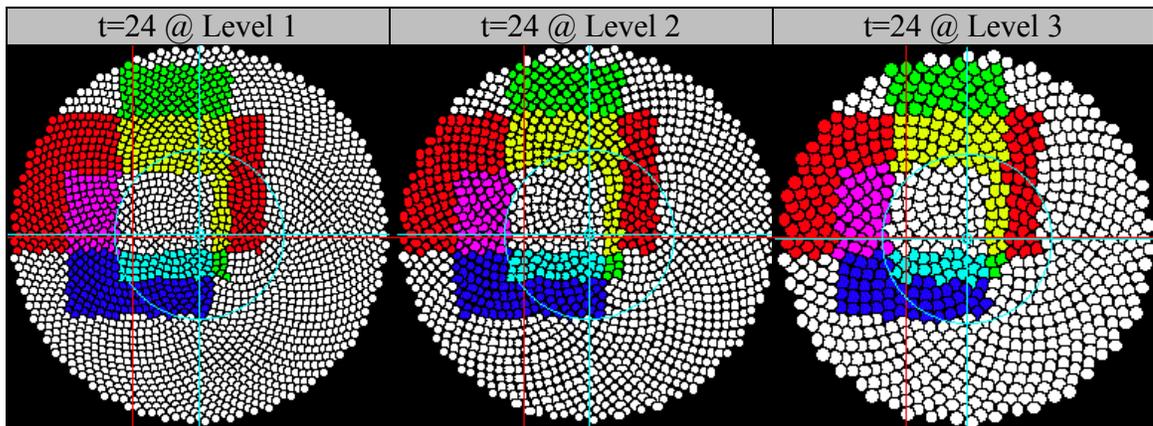
Illustration 257: Imagination filling-in details at tick 24 for all RF levels (Slow+Move)

The borders are considered as straight within the tolerance of a single RF.

D.3.3 Discussion

The imagination filling-in is robust under eye movement, and the visual "thought" is imagined in spatiotopic coordinates (red axes). We wondered if eye movement had any effect on filling-in performance. The
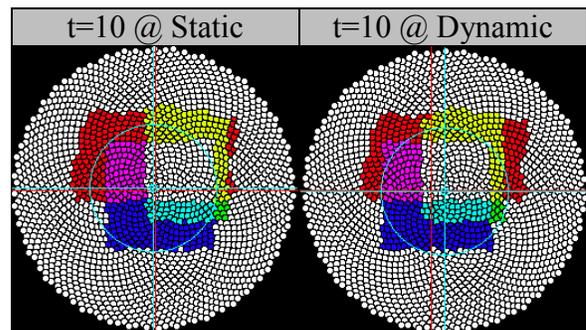
Illustration 258: Imagination filling-in details compared at tick 10 for all RF levels

details of the eye under static and dynamic conditions are compared in Illustration 258. The dynamic eye is moving rightwards, so filling-in on the right is faster (see the red rectangle) while filling-in on the left is slower as compared to the static eye condition.

This could be one more example of the many distortions that occur under eye motion (J. Ross et al., 2001).

D.4 Test 4: Imagine distorted

We wondered how filling-in of imagination would occur under more human and heterogeneous eyes where RFs increase in eccentricity and the ratio of R:G cones is more extreme. Would straight lines still appear straight? Would green still appear as green?

D.4.1 Methods

The methods used are identical to Test 2 except as detailed herein.

D.4.1.1 Subject analysis summary

The agent was similar to that used in Test 2-3 except with high eccentricity dependent RF sizes, and a central foveola region with a 10:1 Red:Green cone ratio. The participant's retinal mosaic is shown in Illustration 259, while details can be found in the supplement (s427). Note
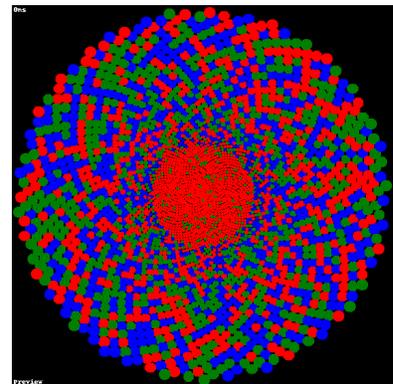
Illustration 259: Imagination Test 4 (Distorted) agent

that imagination occurs within the RFs and not within the retina.

D.4.2 Results

The visual frame-by-frame response results, with precise timing information, are detailed in the supplementary material (s458). In this section, we summarize and discuss various highlights.

D.4.2.1 Imagination filling-in behaviour

The model's behaviour of filling-in for imagination is summarized in Illustration 260.
Eye jitter starts at tick 7 and is indicated via a grey title background. While every tick is
nominally 10ms, we have not calibrated it with biology. At this point in time, we are only
concerned with relative dynamics.
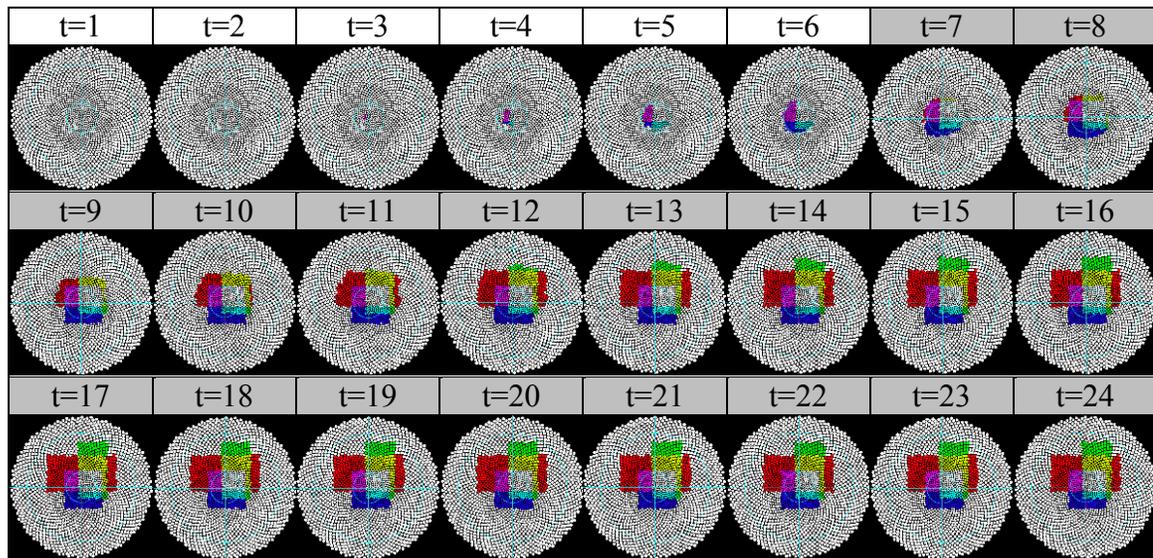
Illustration 260: Imagination filling-in behaviour of 1st Level RFs (Distorted)

Similar spatiotemporal filling-in behaviour occurs at the 2^{nd} and 3^{rd} RF levels. These
details are available in the supplemental material.

Imagination filling-in completed by tick 16. From the perspective of the 1^{st} level RFs, this
took 13 ticks to occur (from tick 3 to tick 16).

D.4.2.2 Imagination filling-in details

The detailed results for all three RF levels are shown in Illustration 261.
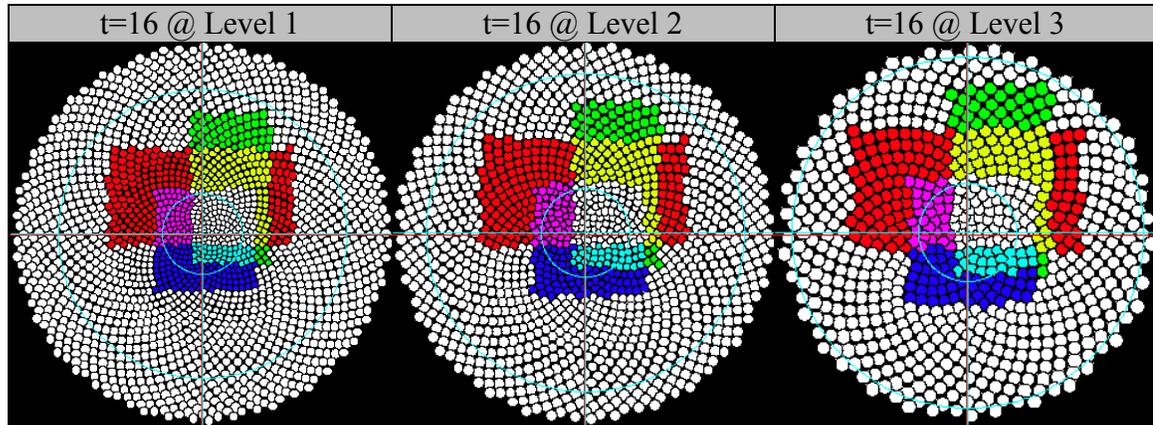
Illustration 261: Imagination filling-in details at tick 16 for all RF levels (Distorted)

The borders are considered as straight within the tolerance of a single RF.

D.4.3 Discussion

Filling-in for imagination times are similar between a highly heterogenic eye than one less so. This may be due to the fact that as one goes up the levels, the RFs become more homogeneous in size, and lateral filling-in occurs at all levels.

Lines are just as straight under the heterogenic eye as compared to the more regular eye. This is because the filling-in coordinates are spatiotopic which is invariant to RF details.

Colours are just as accurate over a foveola with a high Red:Green cone ratio. That is because RFs contain all three L, M and S values and are not dependent on photoreceptors for filling-in. They are dependent on photoreceptors for bottom-up vision, as some RFs may have patches of photoreceptors missing one or two cone types but even these can be eventually filled-in (p9).

D.5   General results and discussion

Top-down imagination is spatiotopic and invariant to both photoreceptor and RF geometries. Straight edges are possible without edge related processing due to how spatial information is encoded.

D.6   Conclusion

We have demonstrated a model for filling-in from imagination. While the initial visual "thought" may progress as a top-down information flow, lateral interactions at all levels of the processing hierarchy are responsible for constructing the sensory image across space. This model demonstrates the power of representing a large and complex image within simple encodings. Of course, the inverse problem of converting an actual image into a simpler object based representation is for future investigation. However, some of the complexities such as the appearance of secondary colours may be due to our human abilities rather than that of the model. This is not a model of conscious perception or conscious imagination – it is strictly sensory.

Our model represents information flows in spatiotopic coordinates so is invariant to eye motion and to the heterogeneities of photoreceptors and receptive fields, either in size or placement. In particular, all our models have emergent border effects, in this case related to the generation of straight lines. There are no edge or orientation mechanisms in the model. That is not to say that future models would not have such mechanisms as these are known to occur in biology. It is just that many mechanistic explanations do not warrant such complexities. The computational functions currently attributed towards much of neurobiology may be incorrect.

We have characterized some of the parameters that may affect imagination. For example, lateral fan-out may influence imagination times. But the point of this toy model is more to get a feel for recurrent processing by playing with it so that recurrence can be more properly incorporated into future plausible architectures.

Appendix E: Solution Overview

Our computational solution consists of the following three major components:

The Emergic Approach forms the inductive foundation for our research as shown in Illustration 262. It embodies *how-to* knowledge for decomposing the complex dynamic system of brain and mind into Emergic Metaphysical concepts – a calculus on a system of equations. In this thesis we focus primarily on the end
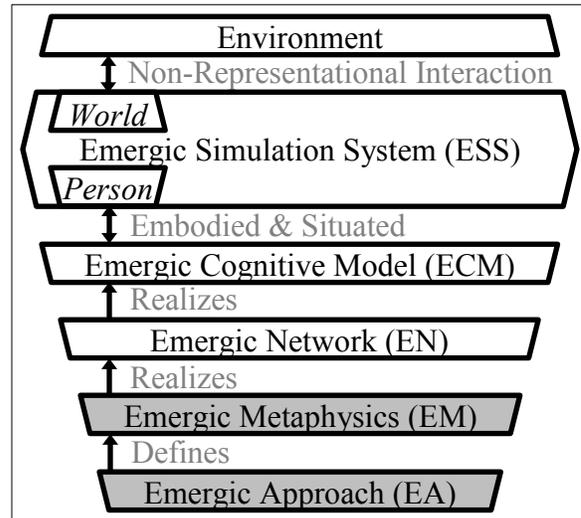
Illustration 262: Solution Overview

product – the computational solution – so the explication of the emergic approach as well as the metaphysics is relegated to future development and exposition.

Our computational solution, therefore, starts with an Emergic Network that realises the concepts behind the Emergic Metaphysics. The Emergic Network is a general simulation system that focuses on the interaction of time delayed flows of information. It can be considered as a hybrid architecture between connectionism and computationalism by cutting up "functionalism" into finer-grained joints to be distributed over space (by harnessing massive recurrence) and over time (by harnessing continuous change), yet

simplifies by using standard computer code to focus on the interaction of information flows. While traditional connectionist units have multiple inputs but only one output, our emergic units are novel in having two or more output flows and can, therefore, directly compute the functional interactions of all information flowing through the unit. It is this focus on massively recurrent interactivity which is known to lead to emergent effects that provides the *emergic* prefix throughout our work.

The Emergic Cognitive Models (ECM) forms our unified cognitive theory for visual filling-in. It is built using the Emergic Network architecture and therefore *is* an emergic network. While the Emergic Network forms the *metaphysical* architecture of ECM, ECM is cognitively decomposed into two emergic functions. The *handle missing data* function and the *maintain information coherence* function form the *cognitive* architecture of ECM from which all behaviours emerge via cognitive re-use.

The Emergic Simulation System embodies and situates a cognitive model in an environment to test out various forms of interaction. The ECM is developed (or grown) inside a virtual *person* which can form visual *views* onto particular *worlds* that generate particular visual stimuli.

We now detail each of these major components in turn.

Appendix F: Emergic Network

The Emergic Network (EN) architecture (Leibovitz & West, 2012a) facilitates cognitive re-use and unified cognitive modelling (Leibovitz & West, 2012b) by providing a principled process metaphysical base – a calculus on a system of equations – on which to found the Emergic Cognitive Model (ECM)
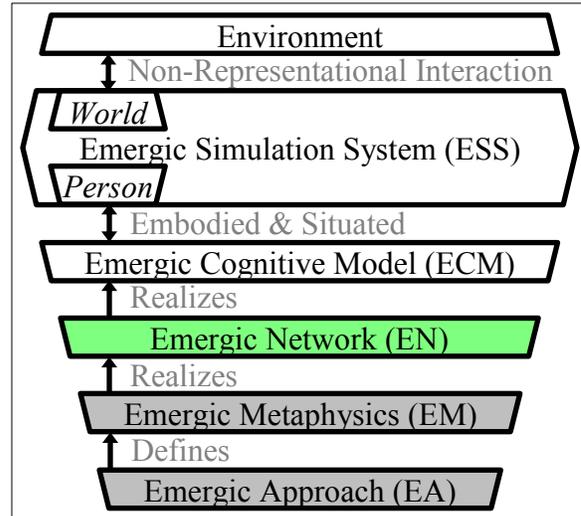
Illustration 263: Emergic Network (EN) Context

demonstrated in this thesis. The context of the Emergic Network is shown within Illustration 263. In this chapter we describe the structure and behaviour of the Emergic Network much like in a reference manual describing a particular solution. Examples of its use can be found for pedagogical purposes (Leibovitz, 2012b) and with ECM (p323). The Emergic Network software is available as open source (Leibovitz, 2012b).

The Emergic Network is a hybrid architecture that takes the best of connectionism and computationalism. From connectionism, Emergic Networks distributes fine-grained "functionalism" over space (via massive recurrence) and over time (via continuous change). From computationalism, Emergic Network uses standard computer code to focus on the interaction of information flows. This focus permits the simplification of classical network architecture into a unit's code.

An Emergic Network could be thought of as a digital version of a recurrently delayed coupled neural network (RDCNN) where the local coupling interacts with no delay (that's our Emergic Unit), while long range couplings exhibit signal propagation delays (Song, Makarov, & Velarde, 2009).

F.1 Entities

An *Emergic Network* contains a set of *Emergic Units* as shown in Illustration 264. Each unit computes an interaction of flows. The flows are segregated into a unit's input or
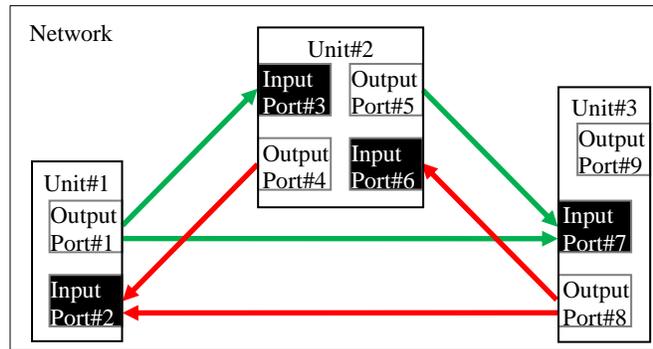
Illustration 264: Emergic Network (EN) overview

output *Emergic Ports*. The output port of one unit is connected to the input port of another unit via time delayed *Emergic Links*. Many-to-one and one-to-many configurations are possible. Once fully specified, including the computations performed, this forms the static structure or architecture of an Emergic Network. The only dynamics in the network is the continual flow of structured *Emergic Values* across links as prescribed by unit computations.

In Emergic Metaphysics, the structure or statics of the system defines the interactive dynamics of the system. However, it is only these interactive dynamics that have causal powers. Causality is not of the form where dynamic *A* causes dynamic *B*, but of the form where dynamic *A* interacts with dynamic *B*. In an Emergic Network, the interactive dynamic components are Emergic Values that are considered to continuously change as

they flow throughout the system. The information flows interact at the locus of an Emergic Unit where the predefined interactive changes are ensured.

In other words, it is the Emergic Values that are causal and the system is best understood when functionally decomposed across *flowcentric* or *infocentric* lines, than across structural or *neurocentric* lines of analysis. That will become a significant finding in the thesis when applied to a cognitive model.

F.1.1   Emergic Network (entity of EN)

The Emergic Network entity represents the entire portion of the brain and/or mind being modeled. The network contains the *entire* set of computational Emergic Units linked together in a highly recurrent manner.

The network represents the *internal* component of simulation. It is embodied and situated within *Person* via the Emergic Simulation System to enable non-representational interaction with an *external* environment.

F.1.2   Emergic Unit

An Emergic Unit computes the interaction of several flows of information segregated by Emergic Ports. In Illustration 265, the input (●) and output (○) ports are shown only for the V1 unit – likely representing a spatial field of units such as receptive fields. In detailed diagrams such as Illustration 264, ports would be expanded, named, and the computation specified.
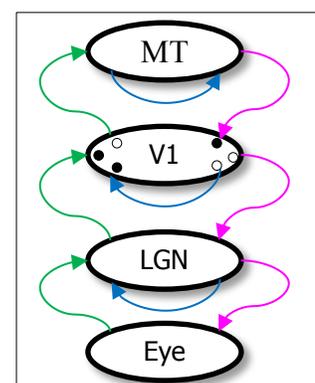
Illustration 265: Emergic network example for vision

318

Physically, a unit can scale from representing a single sensory or motor neuron with minimal computation, to large neural assemblies (Buzsáki, 2010; Grinvald, Arieli, Tsodyks, & Kenet, 2003) with more complicated computations. However, the computations must not be overly complex to retain plausibility. In particular, units do not maintain states, so each is an example of pure functional programming. States (e.g., memories) can be maintained across time by sending a value laterally back to the originating unit. However, lateral flows can have other functionality as well. In "3 Chapter:" (p88) they shift information to another unit in order to maintain visual stability under incessant eye motion. In these cases, the memories-in-motion are best thought of as belonging to the flow rather than to the unit.

The unit's computation is considered to occur instantaneously and continuously so that the conceptual semantics of flow are retained. However, Emergic Links effect some temporal delays to represent processing time.

F.1.3 Emergic Port

An Emergic Port identifies one flow from another within an Emergic Unit. Two flows are colour coded in Illustration 264. Thus, the connectivity among units is actually via their ports. An input port (●) can only receive information, while an output port (○) can only transmit information. The ports are only indicated for the V1 unit of Illustration 265 as this diagram doesn't warrant such a level of detail. When ports are identified and named, the formulae describing their causal effect would also be specified.

In a one-to-many linked output port, the port would duplicate its value to every link. In a many-to-one linked input port, the port could adopt several strategies. It can keep the first, last, or maximum value for example. However, the default is to sum the values received into a distribution.

F.1.4   Emergic Link

An Emergic Link directs the dynamic flow of Emergic Values from the output port of an Emergic Unit, to the input port of another unit (or possibly itself). Both one-to-many and many-to-one configurations are allowed. A link effects a temporal delay and one challenge of designing Emergic models is in accounting for these delays, especially when implementing algorithms originally designed with no delays, or in bringing information together from multiple pathways of differing delays.

Thus the state of the system consists in all the values held within ports and links. Physically, a link can be considered as a set of dendrites, axons or paths among gap junctions (Bloomfield & Völgyi, 2009).

A link can perform minor changes to the values in transit. For example, it can multiply all the values by a static scaling factor (which defaults to 1). Such weights are engineered and do not change unlike classical connectionism. In an Emergic Network, learning must be due to an explicit structure and involves the online dynamics of information flows.

F.1.5 Emergic value

An Emergic Value conceptually represents a continuously changing piece of information that continuously flows across time-delayed links. Computationally, a value is sent as a message, but this does not have to be continuous. For example, rather than sending a zero, nothing needs to be sent at all. This is exactly how neural activation potentials/spikes are delivered or not. An input port can discern whether it has received any value. If none of the ports have received a value, the unit will not be scheduled for computation.

The informational bearing content of a value can be as simple or complex as required to shift the analytical burden away from encodings, and onto the interactions among flows. In ECM, a value in the visual sensation pathways represents sampled distributions (p352), while it represents shift information in the motor planning pathway (p350).

F.2 Software Scheduling Algorithm

The Emergic Network operates with all time in terms of integral Δt tick units from initialization. It maintains one queue for scheduling values to their destination port arrival times, and a second queue for scheduling units for their computation time. The scheduler loops over the following two steps infinitely, and for the current time tick

1. Dequeue all values that are scheduled to arrive now at any destination port, and pass them in an arbitrary order to their destination ports for minimal processing. Summing ports would simply add the new value to the running total. Summing is an abstract method as values are of abstract structure. If a many-to-one destined

port receives multiple values, then if it has semantics such as to take the first or last, then the operation may not be well defined over time.

Then enqueue this unit for computation at the current time tick. It is possible to have a non-zero *input delay* such that once the first port receives a value, this delay allows the port, and others in the unit, to continue to receive (and possibly sum) values for the input delay time, i.e., the input delay effects temporal summation.

2. Dequeue all units scheduled for computation now and invoke their abstract computation method in arbitrary order. In this case, order does not matter as there are neither global nor shared parameters – all the values are available within the unit's own ports.

As part of computation, a unit can check whether a port has a value or not. Typically it sends new values to its output ports. This represents the computation of flow interactions. As every link has a minimal *output delay* of 1 tick, infinite loops within zero time are prevented. An output port would duplicate the value for one-to-many links, and enqueue each value to their destination port depending on the output delay of the link.

Appendix G: Emergic Cognitive Model (ECM)

The Emergic Cognitive Model (ECM) represents our unified cognitive theory of visual filling-in. ECM (Illustration 266) is a single Emergic Network that is "grown" and embodied within a *person*, then situated by the Emergic Simulation System so that it can undergo non-representational interaction with some environment or *world*.
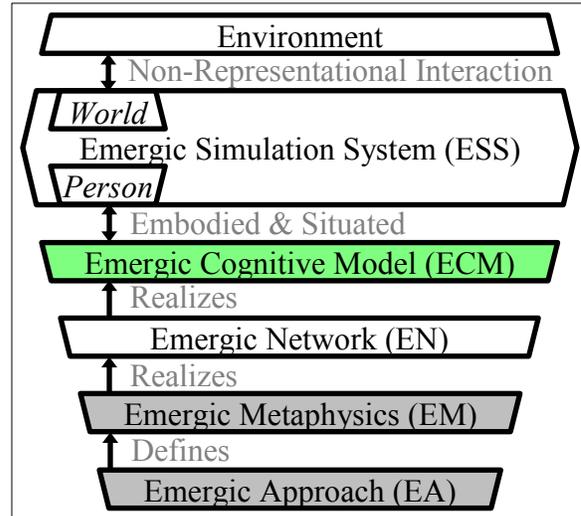
Illustration 266: Emergic Cognitive Model (ECM) Context

G.1   Overview

Unlike engineering, biological functions, especially of the cognitive variety, can be distributed across the forms of the architecture, and every form is a locus for multiple functions where dynamic tradeoffs occur. It therefore does not make sense to ask for the function of a
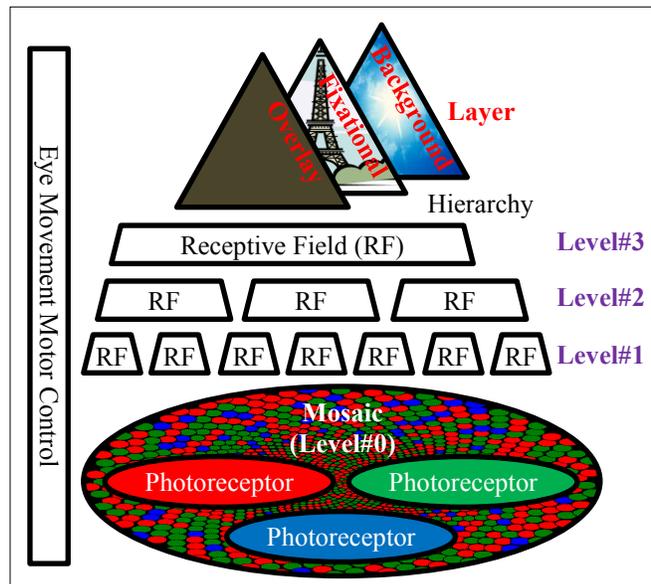
Illustration 267: ECM structural overview

form as it subserves many. Indeed, under *flowcentric* thinking, many functions are properties of the flow and not of the structure. In any case, a given function might not be understood without all the related forms. Thus a holistic overview of the system is required before the details along functional or structural lines can be specified. In this

section we summarize the two high level functions of ECM and the forms of ECM (Illustration 267) over which they are distributed.

G.1.1   Functional Summary

There are two senses to *functions* within Emergic Metaphysics which we distinguish as bottom-up (buFunctions) and top-down (tdFunctions). The buFunctions have a mathematical or computing sense and are represented by a

Illustration 268 Emergic functional decomposition via a unified model

finite set of algorithms (F1 and F2…) distributed over the *finite* structure of an Emergic Network as shown in Illustration 268. They express our *ontological* commitments to cognition precisely at the micro-level of analysis. However, these can be difficult to understand without the context of a higher level.

All macro-level behaviour is completely specified in a bottom-up manner by running the entire Emergic Network. However, because of recurrence, the behaviour of the system (or any one part) in general is not mathematically analytic (Weisberg, 2006) meaning that an *infinite* number of terms would be required to describe it precisely – the apparent complexity of the Mandelbrot set makes a good example (Pollack, 1993). Any *finite* collection of terms would be arbitrary and approximate - a set of epicycles serving a *descriptive* or *epistemic* fitting function. We call macro-level descriptions *epistemic* because they make good heuristics for understanding, illustrating or decomposing the system.

Our tdFuntions are even more approximate in that they do not describe interactions with other tdFunctions (the buFunctions explicitly encode all interactions). For example "too many cooks spoil the broth" and "two heads are better than one" represent dueling epistemic heuristics, but that leaves unsaid how many people should be required for any particular situation. Moreover, while buFunctions compute numeric outputs, tdFunctions are expressed in a form that signifies continuous behaviour rather than computation, e.g., *maintaining coherence*. Indeed, tdFunctions are more akin to biological functions such as *circulation* – not a computation per-se.

Our buFunctions are distributed across the forms and every form subserves multiple tdFunctions. In general, the buFunctions do not help to understand the system or its parts. The tdFunctions are epistemic goals, purposes, strategies, etc. that help to decompose or understand the system. Being epistemic, they have no reality (in code), but great utility for comprehension in setting up context. They therefore appear only within documentation or illustrations. They form a map of sorts. We will describe these first.

The highest tdFunction is that all aspects of a life form, including visual sensation, simply subserve reproductive success – the evolutionary imperative. Seeing is not for the purpose of conscious perception but for action. We translate this imperative to the existing neurophysiological structures as follows.

The tdFunction of visuospatial processing is taken to be the visual analog to the cocktail party effect (McDermott, 2009) – to see the gist of a scene (Oliva & Torralba, 2006; Sampanes et al., 2008) and be able to attend to one Emergic Object among others (Whitney & Levi, 2011), despite occlusions and shadows (B. L. Anderson, 2008), so as to accumulate evidence (S. D. Brown, Steyvers, & Wagenmakers, 2009; Fred & Jain, 2005; Hesselmann, Sadaghiani, Friston, & Kleinschmidt, 2010; Ploran et al., 2007) as to its detailed nature for ultimate decision making purposes (Acker, 2008; Balleine, 2007; Cutsuridis, 2010; de Lange, Jensen, & Dehaene, 2010; de Oliveira et al., 2009; Doya, 2008; Green & Heekeren, 2009; Heekeren, Marrett, & Ungerleider, 2008; Ho, Brown, & Serences, 2009; Newell & Bröder, 2008; Purcell et al., 2010; Shea-Brown, Gilzenrat, & Cohen, 2008; Wang, 2008; Wunderlich, Rangel, & O'Doherty, 2009). As such, a scene is continuously segmented into the background, fixational and overlay layers dependent on attentional shifts. How this high-level tdFunction is distributed among parts will follow.

G.1.1.1   Emergic Object

An Emergic Object is an epistemic and tdFunctional conception (O'Herron & Von der Heydt, 2011; Wilensky & Resnick, 1999). While perception is considered as object oriented (Ciaramitaro, Mitchell, Stoner, Reynolds, & Boynton, 2011; Pylyshyn, 2001), ECM does not have representation for now (Ryan & Villate, 2009; van Zoest, Hunt, & Kingstone, 2010), and possibly never will[2] – therefore no *crisp* objects or *crisp* gist-like

---

[2] Surprisingly we are compatible with a micro-symbolic stance that allows macro-symbols to be created, change over time, and to effect change elsewhere. This turns representation into causally active information flows that are not centrally interpreted. Representational approaches conflate non-causal symbols used for communication and causal "symbols" used for cognition.

proto-objects (Orabona, Metta, & Sandini, 2007). Instead, we simply *consider* all the

information distributed across all units that happen to be associated with some external

object as an Emergic Object. It is the virtual agent (ECM), however, that would decide

what features are used to segment an object (Kaul & White, 2010; Ling, Liu, & Carrasco,

2009; Liu, Larsson, & Carrasco, 2007; Liu & Mance, 2011; Liu, Stevens, & Carrasco,

2007; Saenz, Buracas, & Boynton, 2002; Serences & Boynton, 2007; White & Carrasco,

2011; Zhou & Desimone, 2011), and so objects are not to be found in the environment

(Jones & Love, 2011) or spotlight (Driver & Baylis, 1989; Fisher, 2006; Kastner &

McMains, 2007; Treue & Martinez-Trujillo, 2007), but are agent relative, e.g., one can

focus on one spectator at a stadium, the "wave" of an audience, or a shadow. The

observed speed of a visual object is also agent relative (see Illustration 280 on page 365)

and depends on the object's speed in the world and the speed of the agent's eye

movement across a view of the world (Ryan & Villate, 2009; Sokolov & Pavlova, 2006).

G.1.1.2   Maintain information coherence tdFunction

Ordinarily, an Emergic Network has massive recurrence so that old observed information

from an Emergic Object might eventually combine (top-down) with new observed

information (bottom-up), possibly for a purpose such as filling-in. However, if the eyes

move, then the old information may become incoherent and refer to incompatible

spatiotopic coordinates – the eye is retinotopic. The purpose of the *maintain information*

coherence tdFunction is to shift information (Burr & Morrone, 2011; Hall & Colby,

2011; Keith et al., 2010; Klier & Angelaki, 2008; Mathôt & Theeuwes, 2010; N. A. Parks

& Corballis, 2010; Rolfs, Jonikaitis, Deubel, & Cavanagh, 2011) to compensate for eye

motion as shown in Illustration 5 (p7). Currently, we only compensate for agent directed

eye movement (Bechtel & Abrahamsen, 2010; Meeter, Van Der Stigchel, & Theeuwes, 2010), but eventually we will account for individual object motion (Baccus, Olveczky, Manu, & Meister, 2008; Katsuyama, Usui, Nose, & Taira, 2011; Lorenceau, 2010; Masson, Montagnini, & Ilg, 2010). This continual shifting of information renders the system *flowcentric*, with the flow having spatiotopic coordinates, while the neurons have delayed retinotopic coordinates.

Motion must be planned in advance and distributed to all spatially aware emergic units so that spatially tagged information (Knutsen & Ahissar, 2009) can be appropriately shifted (Klier & Angelaki, 2008) when sent up, down or laterally. This ensures all confluences of information at a particular emergic unit happen to cohere in terms of external (non-representational) stimuli. If we support movement by up to 5 locally connected receptive field diameters, then downward and lateral links must overlap by that amount. Values are distributed down and laterally to all overlapping regions, but only accepted by the few ports that happen to be at the correct place, possibly weighing the value appropriately depending on overlap. This turns the hierarchy into a routing system (Olshausen, Anderson, & Van Essen, 1993; Zylberberg, Fernández Slezak, Roelfsema, Dehaene, & Sigman, 2010) or local area network (LAN).

G.1.1.3 Handle missing data tdFunction

There are many scenarios that can lead to missing data. For example, photoreceptors are heterogeneous and some sensitivity types could be missing in a region, or RFs can be beyond the reach of the photoreceptors; information can be obscured during an eye blink; and information can be segregated to another layer where it is effectively masked from

the fixational layer. In all these case, to support a minimal amount of filling-in (De Weerd, 2006), missing data needs to be handled. As the normal path for sensory data is bottom up, if that is missing, then use lateral data if possible, and if that too is missing, use top-down data. That is the purpose of the *handle missing data* tdFunction of ECM.

G.1.2 Structure Summary

Human vision is considered to have two separate visual pathways for perception and action (Goodale & Milner, 1992). As shown in Illustration 267, ECM ends up having two similar divisions, but their current rationale serves slightly different tdFunctional purposes. As a minor difference, ECM is a model of sensation and not of conscious perception. However, the only reason for having a motor planning system is to enable the *maintain information coherence* function – it is not currently for action per-se.

The motor control units plan and manage eye movement, while the photoreceptor units sense the relevant portion of their view. The Emergic Simulation System manages the non-representational interaction (Garzón, 2008; Pessoa et al., 1998; Stepp, Chemero, & Turvey, 2011; van Gelder, 1998) of these sensors and effectors with the environment. The photoreceptors are arranged in a mosaic which can be considered as the 0^{th} level of an information processing hierarchy (Albright, Jessell, Kandel, & Posner, 2001; Breitmeyer, Öğmen, & Chen, 2004; Spratling, 2002) – information flows not only up, but interacts with other flows going down and laterally (Perlovsky, 2007). Higher levels of the hierarchy are constituted by receptive field units (Fitzpatrick, 2000; Lennie, 2003; Womelsdorf et al., 2008) of increasing spatiotemporal scales (Barrett & Seth, 2011; Horwitz & Husain, 2007; Kayser, Montemurro, Logothetis, & Panzeri, 2009; Kevrekidis

& Samaey, 2009; Kourtzi & Huberle, 2005; Roopun et al., 2008; Rosenbloom, 2011; Silberstein, 2011; Tononi, 2008; W. H. Warren, 2006; Wischnewski, Belardinelli, Schneider, & Steil, 2010). This forms one dimension of the hierarchy. The hierarchy is also divided into a dimension supporting three image segmentation layers – overlay, fixational and background. This segmentation is only minimally implemented. Currently, the model can be considered as realizing the fixational layer. However, in many of the experiments or tests, it can be arranged (without code) for dark information that would normally be perceived during eye blinks or upon occlusion, to be segregated into the background (or overlay) layers where it is effectively masked.

G.1.2.1   Photoreceptor Emergic Units

Each virtual eyeball is looking at a particular view of an environment or world represented by pixels on a virtual computer monitor. The emergic cognitive system transfers this visual information into each computational photoreceptor emergic unit as in Illustration 269 via simple geometry (Illustration 270), and this forms the start of the

Illustration 269: Photoreceptor-Pixel Interaction
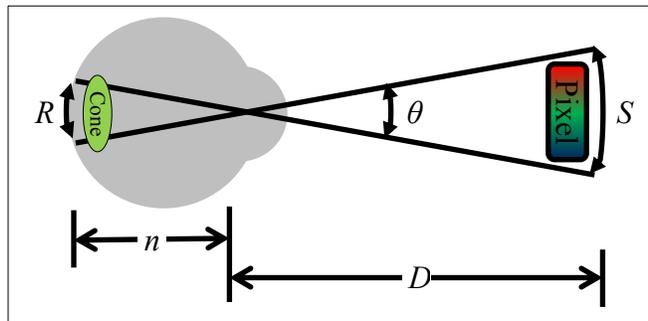
Illustration 270: Pixel to cone conversion

cognitive visuospatial sensory processing. An ECM photoreceptor cuts nature in slightly different joints and models a little more than a biological photoreceptor – it collapses

much of the retinal processing as shown in Illustration 271. The photoreceptors are

arranged in a mosaic that is described next.

The photoreceptors transduce the RGB pixelated view (Luo, Kefalov, & Yau, 2008) into LMS colour space (Atvars & Fomins, 2009; Brill & Robertson, 2007; Pouli, Cunningham, & Reinhard, 2011; Wuerger, Atkinson, & Cropper, 2005) and send this information up the receptive field hierarchy. There is no *lateral* colour opponent processing (Dacey, 1999) in ECM, but eventually there will be top-down and bottom-up *vertical*
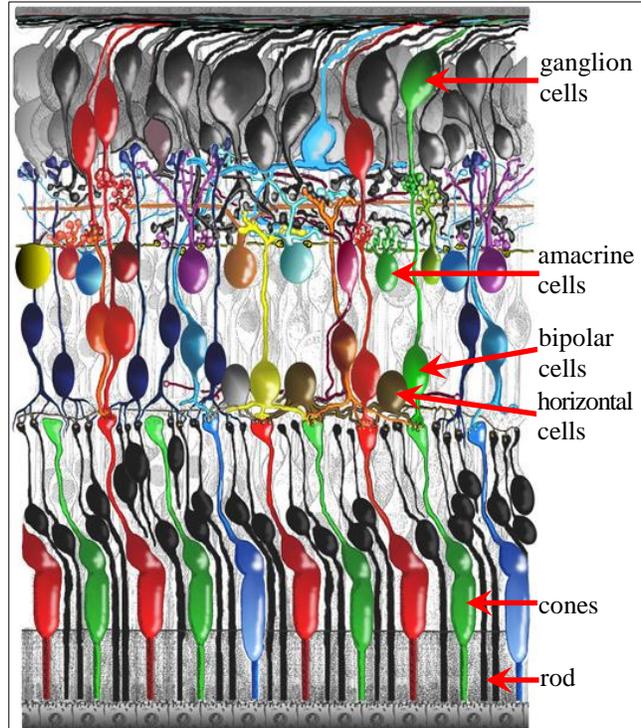
Illustration 271: Biological photoreceptors in retina
Adapted from (Kolb, 2005)

opponency to explain complementary colours rather than opponent colours (Livitz, Yazdanbakhsh, Eskew Jr., & Mingolla, 2010; Pridmore, 2008a; Roelofs, 1961; M. H. Wilson & Brocklebank, 1955). Additionally, photoreceptors detect a change in their value which is also sent up the receptive field hierarchy as an additional flow. This change can be used as temporal edge detection, but it is not yet part of our cognitive theory. It can be considered as temporal opponency. Temporal edge detection is illustrated for all tests within this thesis as an extra **Photoreceptor Changes** column in the animated results online (Leibovitz, 2012a).

G.1.2.2 Photoreceptor Mosaic

The photoreceptor mosaic (Ahnelt, 1998) can be considered as the 0^{th} level of the receptive field hierarchy described next. However, while each of our receptive field units is generic, photoreceptors come in various L, M or S types, and their arrangement has significance to the analytical space.
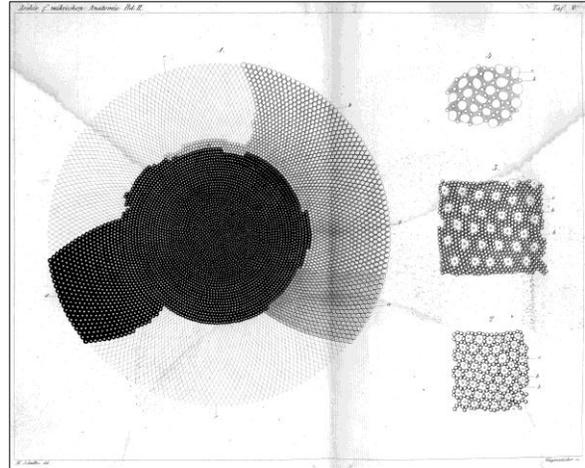
Illustration 272: Retinal mosaic
Drawn by (Schultze, 1866)

The mosaic (as a computational structure) does not participate in the Emergic Network, but it is responsible for developing each photoreceptor emergic units as part of embodying ECM. In particular, the individual photoreceptor are placed in a sunflower pattern (Vogel, 1979) as shown in Illustration 272 with a
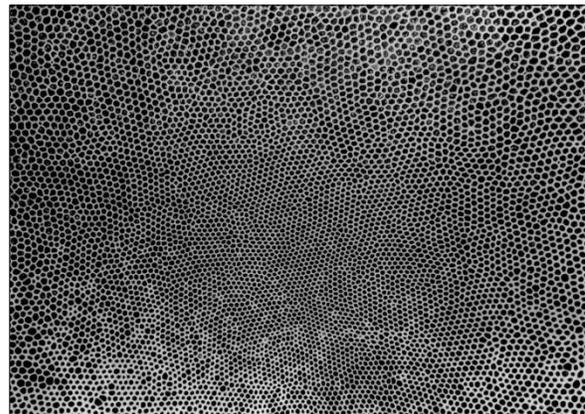
Illustration 273: Human fovea
Tangential section through cone inner segments (Ahnelt, Kolb, & Pflug, 1987)

diameter increasing with eccentricity (J. R. Wilson & Sherman, 1976). Additionally, this mosaic is divided into two zones, the foveola that only has L and M (Red & Green) colour cones (Magnussen et al., 2001) and the rest of the fovea as shown in Illustration 273. The "colours" of the individual cones are random.

333

G.1.2.3 Receptive Field Emergic Units

The Receptive Field (RF) is a computational Emergic Unit and intended to be arranged within a hierarchy of differing spatiotemporal levels above the photoreceptor mosaic. A receptive field simply serves as a spatiotemporal integrator of the values it senses on the way up. These would be the photoreceptor colour space values, as well as their detected changes. While a cone is one of L, M or S colour sensitive types, a receptive field can contain all three which it keeps segregated as pathways, streams (Westheimer, 2007) or simply flows. Receptive fields will send their values up for further spatiotemporal summation (R. A. Smith & Cass, 1989). They will send them laterally to support image stability (Martinez-Conde et al., 2008), Emergic Objects and filling-in (Anstis, 2010). They will send information downwards also to support filling-in. Missing colour values are filled-in from lateral or top-down flows, and could be linearly interpolated if required, although this is turned off for all tests in this thesis. Some integration occurs automatically via summing input ports, and with colour values that represent statistical distributions (Alvarez & Oliva, 2009; W. J. Freeman, 2007; McAlpine & Palmer, 2006; Panzeri, Brunel, Logothetis, & Kayser, 2010) supporting integration and Emergic Object functionalism. In particular, these values are explicit, can be structured, and remove the encoding strategy (Spratling, 2004) level of analysis from Emergic Networks to focus on massive functional interconnectivity instead (Hohwy, 2007).

G.1.2.4 Receptive Field Level

A receptive field level (as a computational structure) is similar to a photoreceptor mosaic. It develops a similar sunflower pattern of receptive fields, although each level contains larger receptive fields (J. Freeman & Simoncelli, 2011). It also does not participate in the

Emergic Network, but serves to develop the receptive field emergic units and link them to lower receptive fields or photoreceptors as required.

G.1.2.5 Receptive Field Hierarchy

The receptive field hierarchy (as a computational structure) does not form part of the Emergic Network but develops the receptive field levels (bottom-up) and photoreceptor mosaic – all within an embodied hemifield (Large, Culham, Kuchinad, Aldcroft, & Vilis, 2008) of an Emergic Simulation System *Person*. Only the photoreceptors and RFs form part of the emergic network.

While the receptive field levels form a hierarchy, their overall spatial extent can be larger than what all the photoreceptors could sense at any one time (Hubbard et al., 2010; Intraub, 2012; Park et al., 2007). While this is filled-in generically, it also serves as an imagination and knowledge buffer – to know what is outside the scope of current vision.

Photoreceptors transmit individual colour values upwards to the summing ports of first level receptive fields so that they can be spatiotemporally integrated. Receptive fields transmit each colour value to the summing ports of higher level receptive fields to support both spatiotemporal integration as well as Emergic Objects. These values can be sent down and laterally to support filling-in.

G.1.2.6 Segmentation Layer

A visual scene is continuously segmented (Eckhorn, 1999; Marcus & Van Essen, 2002; Scholte, Jolij, Fahrenfort, & Lamme, 2008) into overlay (B. L. Anderson, 2008),

fixational and background layers (Baijal, 2009; Fang, Boyaci, & Kersten, 2009; Lawlor, Holtmann-Rice, Huggins, Ben-Shahar, & Zucker, 2009) dependent on attentional shifts (J. M. Brown, 2009; Buschman & Miller, 2010; Compte & Wang, 2006; Drummond & Shomstein, 2010; Herrington & Assad, 2009, 2010; Ikkai & Curtis, 2008; Kelley, Serences, Giesbrecht, & Yantis, 2008; Paffen & Van der Stigchel, 2010; Womelsdorf et al., 2008, 2006). While this is future work, it is essential in understanding the overall function of visuospatial cognition. Note that a hierarchical receptive field *level* and a segmentation *layer* are two different dimensions with differing functional purposes.

G.1.2.7   Eye Motion Planning & Distribution Structure

While it is a requirement that both the environment and the agent engage in autonomous motion, e.g., Illustration 280 on page 365, the *advance* planning of eye motion, and the *distribution* of this information to all receptive fields was necessitated by the *maintain information coherence* tdFunction rather than any requirement for motor control or image stability. When information is sent up, it is never shifted by design to maintain maximum fidelity (Frechette et al., 2005). When information is sent laterally, its position is adjusted by the amount this level is about to move so that upon arrival, the confluence (De Kleer & Brown, 1984; Forbus, 2011; Westheimer & Ley, 1997) of information is coherent. When information is sent down, it is adjusted based on the amount this and the lower levels are about to move. Thus movement compensation needs to know motion plans two moves in advance. This is biologically managed by corollary discharge that acts as an advance (efference) motor command copy distributed electrically (Perrone & Krauzlis, 2008). Because the higher levels have large receptive fields, this shifting amounts to updating information within the receptive field, rather than moving information to a

lateral receptive field as would occur in the lower levels. If movement is beyond the lateral connectivity, information is effectively masked and extensive smearing can result (Holcombe, 2009) and page 96.

G.1.3   Structure development and embodiment

The structure of an ECM is generated by a development program, (e.g., Koene et al., 2009), that embodies (M. L. Anderson, 2003; Bridgeman & Tseng, 2011; Kiverstein & Clark, 2009; Sloman, 2009; van de Laar & de Regt, 2008; M. Wilson, 2002) the ECM within a virtual agent, or *person* partly shown in Illustration 289 on page 377. The structure is heavily parameterized (Breakspear, Jirsa, & Deco, 2010; Eliasmith, 2007; Shiffrin, 2010) at a high-level of description providing a space of possible realizations. So while a generic structure can be described, indeed, that is what the development program creates, an actual instance has a stochastic and parameter dependent architecture. It is quite possible that these parameters interact in an emergent fashion.

The development program (Fair et al., 2009) is only used to create an embodied *person* at the start of a simulation. The *person* does not develop over time (Johnson & Munakata, 2005) as that would imply that the structure varies, i.e., it is no longer structural – we try to maintain a principled distinction between structure and dynamics. Moreover, the structural products of classical learned systems can be so complicated and distributed as to be incomprehensible (Domingos, 1999; M. Page, 2000; Parker, 2010; Wagner, Schmuki, Wagner, & Wolstenholme, 2006; Zgurovsky & Pankratova, 2007). Thus, while ongoing development is possible, that is not the focus of our research. ECM incorporates all learning within the dynamics of information flows.

G.1.3.1   Parameters

While the emergic cognitive system can store a *person's* parameters, it cannot store its current state. Thus every time a *person* is loaded into the emergic simulation system or regenerated, its resulting structure will differ in detail and its behaviour (such as filling-in times) may be variable. Where this is critical, multiple tests ought to be run as with human experiments. In this thesis, such multiple runs were not necessary as the tests are used mostly to uncover qualitative behaviour.

Static parameters along with stochastic developmental choice determine individual differences, while dynamic parameters specify a *person's* location in space, where they are looking, etc.

The *person's* parameters are described starting on 378, and, how the development program makes use of these parameters is described next.

G.1.3.2   Development Program/Architecture

The development program generates the static structure or architecture of a *person* based on the parameters that allow for individual stochastic variation. It does not continue beyond initial generation. Each modality integrates portions into the overall development program. Currently, the visuospatial modality is the only one designed.

The generic development program is hierarchical and straightforward. A single Emergic Network is created to house the brain and mind. Two embodied hemifields are created to

house bilateral components, but they are empty in the "core" *person* case. The visuospatial modality does add bilateral Emergic Units associated with each eye.

Thus the static structure of an empty or core *Person* is made up of one Emergic Network and zero Emergic Units, Ports and Links. As a consequence, and without sensory interactions, there are zero dynamic Emergic Values flowing through the system. Note that Illustration 289 on page 377 shows non-zero values for Emergic entities only because it already includes the visuospatial modality.

A core *Person* is filled-up as follows. Including the visuospatial modality, the order of ECM construction is one brain (Emergic Network) and two hemifields (one for each eye). In each hemifield, construction proceeds as one motor plan distribution system, one photoreceptor mosaic, followed by one receptive field hierarchy. The mosaic creates the photoreceptors in a sunflower pattern. The receptive field hierarchy creates three receptive field levels bottom-up, linked to lower levels or the mosaic and the motor plan distribution system. The lateral and downward links cross a much larger spatial extent to support Emergic Objects and the *maintain information coherence* function. The two hemifields are currently not under any common control, so stereoscopic vision is not possible. For this reason, all tests in this thesis involved one eye. Indeed, motion (beyond jitter) is currently under the control of the investigator (or operator) as there is no attentional system planning the next saccade.

G.1.3.3 *Person* vs. cognitive parameters

In Emergic Metaphysics, both software and hardware are structural, but there are times where it is useful to distinguish the two. Similarly, under a Universal Turing Machine (or any modern operating system), a program is just data, but there are times where it is useful to distinguish the two. In the analogy of a word processor, its behaviour, especially over time when reformatting the entire document will differ depending on the amount of hardware memory, and if the document contains much graphics, the performance may also depend on the hardware graphics systems. The point is that these hardware parameters are considered as different from the software parameters of the word processor, even though both might contributed to overall performance behavioural differences while engaged in word processing.

In an Emergic Network such as ECM, the overall design is governed by epistemic tdFunctions which distribute ontological buFunctions across computational units. These represent the unified cognitive model of visual filling-in and these buFunctions are parameter free, i.e., any possible parameter has been fixed and remains unchanged throughout unified testing. However, given the basic design, every *Person* can develop with a different arrangement of units governed by the overall architecture. These developmental differences are governed by numerous parameters and will be explained next. However, the computational system remains homogeneous and with zero parameters. For example, every RF computes exactly like any other RF.

G.1.4 Flows of information (Emergic Values) within major pathways

The only dynamics in an Emergic Network such as ECM is the flow of Emergic Values throughout the system. As the values transit through Emergic Units, they undergo change. In this section we summarize the various causal flows of continuous change within ECM. These flows follow two main pathways.
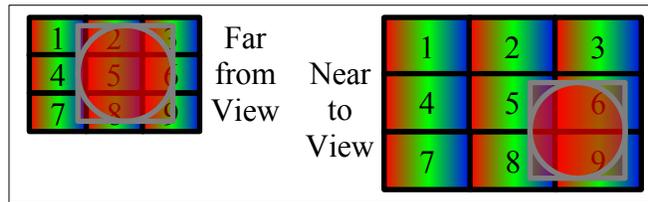
Human vision is considered to have two separate visual pathways for perception and action (Goodale & Milner, 1992). ECM ends up having a similar two pathways, but their current rationale serves slightly different functional purposes. Each has its own set of flows. These pathways interact at the locus of receptive fields.

G.1.4.1 Visual sensation pathway

The visual sensation pathway currently contains only information about surface L, M and S colours at various spatial locations. If information is missing, the *handle missing data*
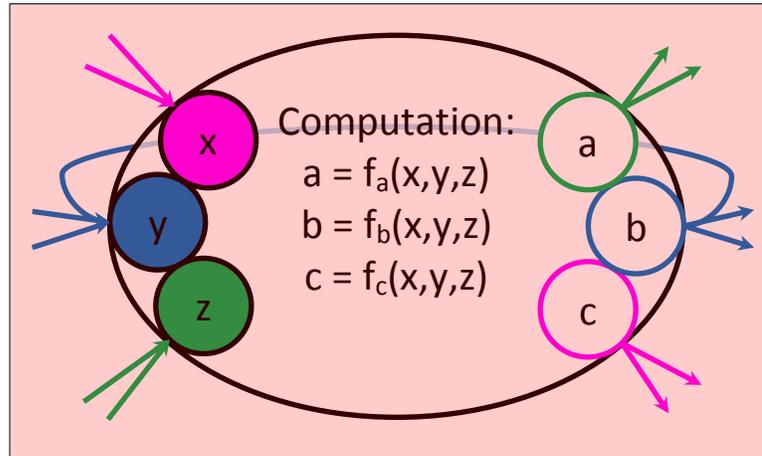
Illustration 274: Possible architecture of three flows

tdFunction will fill-in the data. Information flows in three directions, bottom-up, top-down and lateral.

While a possible three-flow architecture is shown in Illustration 274, the actual Receptive Field architecture is shown in Illustration 275, where it is the connectivity by Emergic Links to the appropriate Emergic Ports that determines if a flow goes upwards, downwards or laterally.
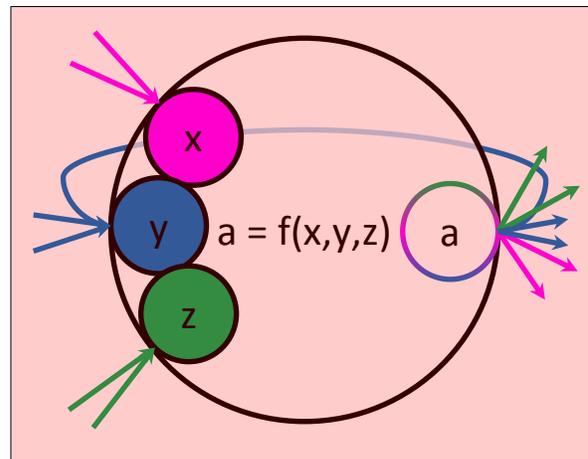
Illustration 275: Actual architecture of three flows

If there was no eye motion, then connectivity would be as follows. Going upwards, there would be a many-to-one connectivity to achieve spatial summation in a perfect hierarchy. Going downwards would be a symmetric one-to-many connectivity so that a missing

value could be filled-in by the spatial average. There would be no lateral connections, but a single recurrent connection to self. This would allow a missing value (such as during a blink) to be filled-in by the previous (temporal) value from before. This is required because Emergic Units cannot store state, but can send a value back to themselves. Also, if there was no eye movement, the buFunction $f(x,y,x)$ would be equivalent to the guiding *handle missing data* tdFunction as the the guiding *maintain information coherence* is essentially nil.

However, there is eye motion, and connectivity becomes more convoluted. The upward path is identical and is engineered to be invariant under eye motion. This was a design choice and has not yet been validated with biology. Lateral connectivity is defined by a lateral fan-out development parameter and allows the RF to shift information to compensate for eye motion. This lateral connectivity includes the recurrent connection to self. Going downwards is another fan-out development parameter again to compensate for eye motion.

Thus while the eye is under motion, the buFunction is unchanged, but the corresponding tdFunction, *handle missing data*, becomes attached to the flow where memory-in-motion effects, dissipation, and border effects emerge.

The lateral connectivity is only for the *maintaining information coherence* tdFunction. ECM does not yet make use of surround suppression which would involve even greater lateral connectivity. Regardless, lateral connectivity is extensive in ECM.

"A surprising, but consistent pattern across all areas of neocortex examined, is that the most of the thousands of synapses formed on the dendritic tree of any neuron come from its neighbouring excitatory neurons located in the same cortical area. Very few synapses are contributed by long distance connections, whether they arise from neurons in subcortical nuclei (principally the thalamus) or other cortical areas." – (Douglas & Martin, 2007)

G.1.4.2 Visual coherence pathway (motor planning)

In order to realize the *maintain information coherence* function, RFs must shift their visual content to compensate for eye motion. However, the RFs need to know how much shifting is going on at the time they sense their visual information. The distribution of such information takes time, so proprioceptive feedback would come too late after visual content. Thus, eye motion, including jitter and microsaccades under involuntary control, needs to be planned in advanced and distributed to all RFs[3]. Certainly there is a specific motor planning unit that plans eye-jitter as well as moving the eye as directed by the *Person* controller on page 377. However, the visual coherence pathway ensures these plans are in advance, distributed down the hierarchy then reflected back upwards so that all RFs receive eye movement information at the same time as visual content allowing them to correctly compensate for eye motion. The reflection is an example of an Emergic Link with a two tick delay (as shown in Illustration 278) due to the fact that shifting information does not descend into the photoreceptors.

---

[3] It is not critical to our model whether involuntary controlled jitter requires *advance* planning. It was simpler to plan all motion in one place. It is certainly possible for jitter control to be separated from other eye motion control.

G.2 Functional decomposition details

In classical computer code, there is a multi-level nesting of functions within functions. Barring temporal delays, the lower level functions can be said to exactly realize the higher ones. It is possible to engineer distributed code that is isomorphic to linear programing. Unfortunately, Mother Nature was not such a strict engineer (Brooks, 1985), and classical functional concepts are problematic. Our tdFunctions are epistemic creations to help make sense of ontological buFunctions that are distributed and intermixed with others over a network of units. Despite being epistemic, the tdFunctions are critical theoretical constructs.

> *"All these studies aim to develop a theoretical framework for visual computation or more generally for the encoding and processing of information by the brain. Mathematicians often say that "a well-posed question is one that is already half-answered". We think the same applies in systems neuroscience. Once we are able to develop the appropriate theoretical concepts, half of the problem of understanding brain function will be solved. Data rules, but theory understands."*
> *– (Pouget, Ringach, & Landy, 2010)*

In this section we provide the detailed mapping between our two tdFunctions and the buFunctions distributed across ECM.

G.2.1 Handle missing data tdFunction

An RF receives visual information in three flows (see Illustration 4 on page 6)

1. Bottom-up from the stimulus (via photoreceptors or lower level RFs),

2. Laterally from nearby RFs at the same level, and

3. Top-down from larger RFs that hold local summary statistics

While all three sources could be combined into the single RF output, ordinarily only the bottom-up source is used. The purpose of the *handle missing data* tdFunction is to decide what to do in the case that bottom-up information is not available.

G.2.1.1   Corresponding buFunction pseudo code

The *handle missing data* tdFunction is realized by a single buFunction that forms the sole *computation* within each RF emergic unit. In pseudo code

```
foreach colourSensitivityType 'c' in (L, M, S)

    if rfInputBottomUp[c] is not null

        rfOutput[c] := rfInputBottomUp[c]

    else if rfInputLateral[c] is not null

        rfOutput[c] := rfInputLateral[c] * 0.75

    else

        rfOutput[c] := rfInputTopDown[c] * 0.10

    if rfOutput[c} is not null

        rfOutput[c].send()

    // If not sent, it is effectively masked
```

Note that such pseudo code is procedural rather than purely functional, and certainly differs from traditional connectionism, although some are extremely flexible (Feldman & Ballard, 1982).

The scaling factors used above (0.75 and 0.10) effectively adjust the weights of the sampled values. This is to reflect the fact that lateral values are most precise spatially, but one tick old, while top-down values are spatial averages and also one tick old.

346

G.2.1.2   Corresponding ECM architecture

The three spatiotemporal levels can be thought of as supporting the *handle missing data* tdFunction, as well as the spatial summation flow going upwards and its reflection downwards. For now, a single lateral connection can be thought of as the previous temporal value of the RF.

G.2.2   Maintain information coherence tdFunction

While the *handle missing data* function is almost one-to-one with its buFunction, the *maintain information coherence* tdFunction is distributed across the architecture and makes use of dendritic processing (e.g., Memmesheimer & Timme, 2012). At a high level, eye motion is planned in advance and this information is distributed to all RFs such that visual content can be shifted to account for eye motion. The details include the following aspects that will be described one at a time.

1. The motor planning system and pathway consists of three kinds of Emergic Units. The Planner, multiple Distributors and the Executer. The Planner determines the eye motion in advance and sends this info, via a dedicated Emergic Value, to a Distributer which distributes the information down until it reaches an Executer which causes the eye motor muscles to saccade. It also reflects the information back upwards. Each RF is also a Distributer and can use its spatiotemporal hierarchy to effect the distribution. However, the 1st level RFs simply reflects the information back up after a two tick delay. The information has to go back upwards in parallel with the visual content where shift compensation actually occurs.

2. In the sensory pathway, all visual content is tagged by it spatiotopic coordinates in a manner that allows for spatiotemporal summation via a dedicated Emergic Value. All Emergic Units and their ports are updated with their new spatiotopic coordinates, and when they receive a value, they will extract the percentage of distribution that overlaps with their own spatial extent.

3. Lateral and downward connectivity is expanded via fan-out parameters to allow for shifting.

G.2.2.1 Motor planning system

The motor planning pathway and system is responsible for planning eye motor commands in advance and distributing this information to all RFs so that they can compensate for eye motion when it eventually occurs. The planning system contains three kinds of Emergic Units, the Planner, the Distributer and the Executer as shown in Illustration 276. This pathway uses a dedicated *Shift* Emergic Value in its downward and upward flows.
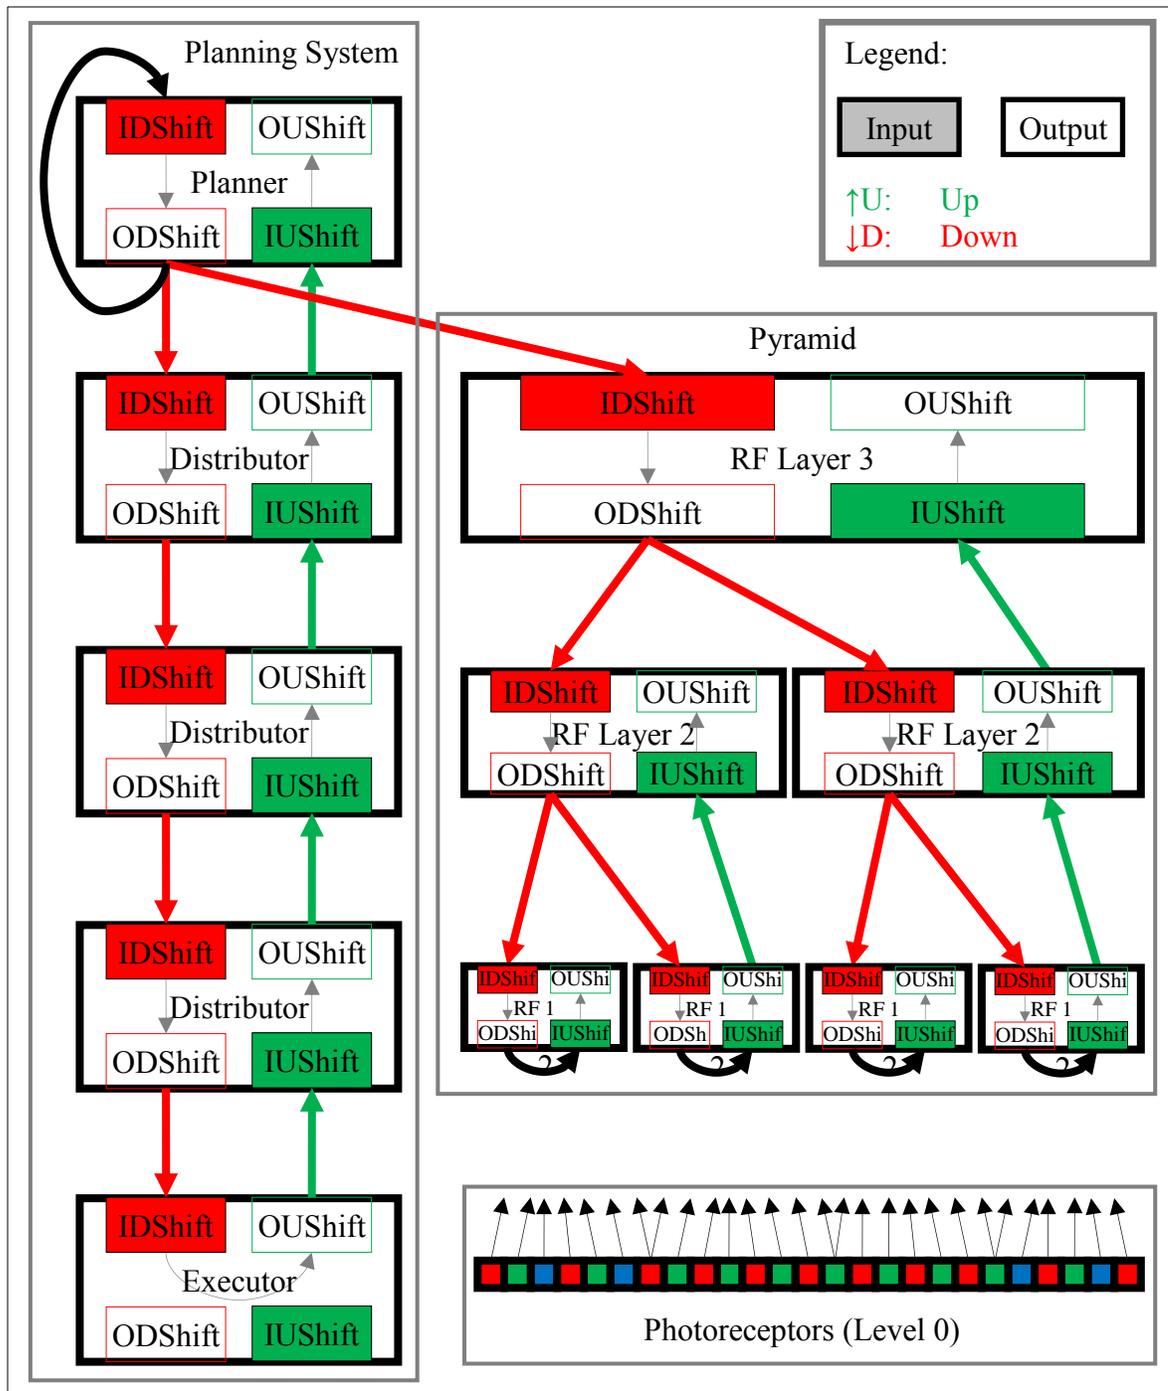
Illustration 276: Motor planning system and pathway

G.2.2.1.1 Shift Emergic Value

The planning system distributes Shift information that has the following Emergic Value structure:

- dX: the change in the eye's X coordinate
- dY: the change in the eye's Y coordinate

These Shift values are not fixational shifts to target, but planned ballistic motor movements broken down into a per tick amount.

G.2.2.1.2 Planner Emergic Unit

The Planner breaks down target fixations into per tick movements. Thus it must send the last partial shift back to itself so it can plan the next one. It adds jitter then sends the Shift information down to its Distributers. This has typically one-to-many connectivity, and the shift information is duplicated to all lower level Distributers automatically. Note that every RF is also a Distributer and manages both the motor planning pathway as well as the visual sensation pathway.

G.2.2.1.3 Distributer Emergic Unit

The Distributer receives shift information downwards and continues this flow downwards. It will eventually receive this information reflected upwards, and will continue that flow upwards. It is this upward path that parallels the visual content and during which compensation for eye motion occurs as shown within Illustration 5 on page 7. Note that every RF is also a Distributer.

The downward path has one-to-many connectivity, and the shift information is duplicated to all lower level Distributers automatically. The 1st Level RFs have a two link delay recurrent loop to reflect the downward flow upwards. The extra delay is to compensate for not sending shift information to the photoreceptors.

The upward path has a sparse one-to-one connectivity as a larger RF needs to receive the reflected shift from at most one lower level RF – all else would be duplicates.

G.2.2.1.4   Executer Emergic Unit

The Executer Emergic Unit effects the eye motor command, and reflects the shift information upwards.

G.2.2.2   Visual sensation pathway

The visual sensation pathway supports surface colour values in a manner that enables spatial summation and the filling-in as required by the *handle missing data* tdFunction. This pathway is extended so that the values are spatiotopically tagged, shifted as required, and are partially accepted by Emergic Ports by the amount of spatiotopic overlap.

G.2.2.2.1 ValueXY Emergic Value

The surface colours are encoded with their spatiotopic extent. This Emergic Value has a structure that is simplest to introduce via its signature:

```
ValueXY((x,y,v,w)…)
```

where

x:	the spatiotopic x-coordinate of this sample value
y:	the spatiotopic y-coordinate of this sample value
v:	the L, M or S colour value of this sample from 0 to 255
w:	the hardcoded weight assigned to this sample value; typically 1
…:	repeatable, in which case this emergic value represents the descriptive statistics for all the samples weighted accordingly

In order for this value to represent descriptive statistics, it is internally structured as follows:

n:	the number of samples $= \sum 1$
sw:	sum of sample weights $= \sum w_n$
swx:	weighted sum of X coordinates $= \sum w_n x_n$
swy:	weighted sum of Y coordinates $= \sum w_n y_n$
swv:	weighted sum of colour values $= \sum w_n v_n$
swxx:	weighted sum of X coordinate squared $= \sum w_n x_n^2$
swxy:	weighted sum of X coordinate multiplied by Y's $= \sum w_n x_n y_n$
swyy:	weighted sum of Y coordinate squared $= \sum w_n y_n^2$
swvv:	weighted sum of colour value squared $= \sum w_n v_n^2$

Because all these internal values are sums, it is trivial to add such values together or to add a sample to a distribution – indeed a sample is nothing more than a distribution with n=1. With such descriptive statistics, it is possible to determine the

mean(x): swx/sw

variance(x): $(swxx - swx^2/sw)/sw$

covariance(x,y): (swxy - swx*swy/sw)/sw

sd(x): $\sqrt{variance}$

for values and spatial extents. In particular, the standard deviations for the X and Y extents are used for eye movement compensation (or "routing" purposes) as will be described shortly.

Weights or scaling factors are hardcoded and typically unitary. Within ECM, lateral and top-down values are effectively reweighed by the scaling factors of 0.10 and 0.75 respectively as shown within the *handle missing data* buFunction pseudo code on page 346.

G.2.2.2.2 Widened connectivity, routing and shifting

To allow for the relative shifting of information to compensate for eye movement, connectivity must be spatially widened. A pure hierarchy becomes physically less so although it remains functionally pure. The system has been engineered so that upward flows do not need to be widened. However, lateral flows have connectivity widened by a lateral fan-out development parameter, and downward flows have connectivity widened by a downward fan-out development parameter. The visual information is essentially broadcast on a widened local network, but only accepted by the ports of RFs that happen

to have the correct spatial extent. Effectively, visual content is shifted as shown within

Illustration 5 on page 7. This makes the visual system act like an routing system

(Olshausen et al., 1993; Zylberberg et al., 2010). If eye movement is greater than the fan-

out parameters, information is effectively supressed, masked and lost. These fanout

parameters affect *person development* but do not affect the *cognitive* buFunctions (and

our model) which remain parameter free (see section G.1.3.3 on page 340 for an extended

discussion).

In Illustration 278, most of the "complicated spaghetti looking" connectivity is due to

widening a simple connectivity to allow for shifting.

G.2.2.2.3   Spatial sensitive Emergic Ports

Without eye motion, a higher-level RF would simply spatiotemporally sum the

information it receives from many lower-level RFs based purely on connectivity as

determined by Emergic Links. However, to compensate for eye motion, connectivity is

greatly increased to allow for the shifting of information. Now, some of the information

needs to be dynamically included, and some is spatiotopically shifted beyond the RF.

This is accomplished by using the RF's spatiotopic extent (as shifted) and comparing it to

the Emergic Value's spatial extent as provided by its spatial statistics.

The standard deviations are used to determine how much of a colour value overlaps with

a port's spatial extent, and whether part of the value is accepted by the port. Ports accept

colour values which overlap their spatiotopic RF boundary by at most one standard

deviation. As an example, the spatial extent of the descriptive statistics within an Emergic Value can be calculated as follows.

mean(x):   swx / sw

variance(x): $(swxx - swx^2/sw)/sw$

sd(x):   sqrt(variance(x))

In this example, a port whose RF is at x and has a radius of r would accept a colour distribution if there is an overlap between x±r and mean(x) ±sd(x). The percentage of overlap determines how much of the sample distribution is taken as part of the port's sum.

An example of how these statistics work on colour values is shown in Illustration 48 on page 57 and in Illustration 277 on page 357.

G.2.2.2.4   Coordinate systems

The visual content is received by the photoreceptors but tagged via their spatiotopic coordinates. Looking only at activation values would make photoreceptors and RFs appear to be retinotopic. However, we claim that spatiotopic information is effectively encoded over the retinotopic spike trains to allow for receptive field remapping (Merriam et al., 2007). The entire RF level has its spatiotopic coordinates shifted to compensate for eye motion as the visual content flows upwards. That implies that during eye motion, while information is flowing upwards, it is not physically shifted, but the effective shift difference between the retinotopic and spatiotopic reference frames increases. It is only when information is sent laterally or downwards that it appears to be shifted according to spatiotopic lines.

The information flow is always spatiotopic. The neurons can act in both a retinotopic and spatiotopic manner depending on how their behaviour is analyzed. During eye movement, a shift or perturbation effectively moves up the hierarchy giving the RFs a delayed retinotopic and spatiotopic coordinate system. This can manifest itself as various kinds of receptive field remappings (e.g., Duhamel, Colby, & Goldberg, 1992; Hall & Colby, 2011).

G.2.2.2.5   Shifting example

There are multiple ways to mechanise and conceptualize the shifting of information that are behaviourally equivalent. Whether the flow is shifted and the RFs are stable, or vice versa may be relatively isomorphic. We had actually implemented both versions and settled on the following due to its simplicity. Biological evidence will be required to arbitrate.

Emergic Values are initially tagged with spatiotopic coordinates that never change. RFs have a retinotopic coordinate that is updated every tick by a Shift amount. RF levels therefore act as if they are transiently spatiotopic (Wurtz, 2008) and continually shifting while the flow of information is stable and invariant to eye motion. RFs act as if they are spatiotopic, and only accept the part of the flow that matches their spatiotopic coordinate. Surprisingly, while an RF acts in a spatiotopic manner, it can be easily misconstrued as acting in a retinotopic manner if the mechanism of encoding is not taken into account.

Illustration 277 demonstrates how the statistical based routing and summing works. In (A) there are four rectangular stimuli patches with identically sized RFs perfectly overlaid (you can see the RF overlays shift in (B)). We name these four RFs after their originally perfectly aligned stimuli grey, red, green and blue. The grey RF has colour information for two R, two
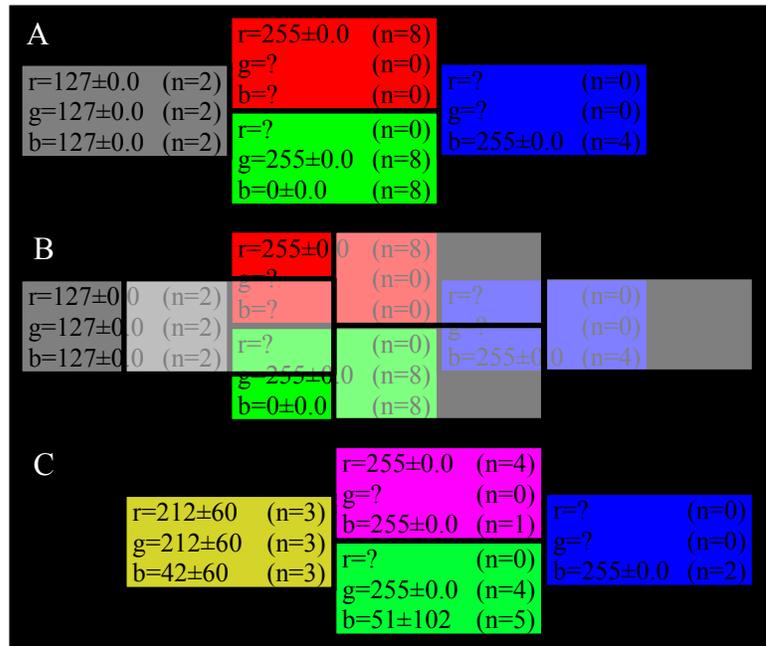
Illustration 277: Statistical shifting
A: Four RFs perfectly aligned with coloured stimuli
B: Eye shifts right by ½ RF width
C: New RF values by statistical lateral shifting

G and two B sensitive photoreceptors; the red RF has information on a patch of eight R only sensitive photoreceptors; the green RF has information for eight G and eight B sensitive photoreceptors; and the blue RF has colour information for a patch of four only B sensitive photoreceptors. We assume the rectangular sized RFs indicate the stimuli's spatial extent and this would be reflected in the standard deviations of the sampled values (p352). In (B), the eye has shifted half an RF width to the right and so there will be a 50% or 25% overlap between appropriate RFs and the old values shifted laterally. We will assume that the eye is currently blinked as well, so the RFs can only retrieve their old values from lateral connectivity. We will further assume the fan-out parameters are 5, so effectively every RF receives lateral connectivity information from all the RFs shown including recurrence from themselves. Nevertheless, the input summing ports to the RF

only accepts the portion of the Emergic Values that overlap, and this is shown in (C). The new grey RF takes ½ its old value, ¼ of the red's old value and ¼ of the green's old value. This produces a greyish yellow. The new red RF takes ½ its old value and ¼ of the blue's old value. This produces a perfect magenta colour despite the difference in RB sample numbers. The new green RF takes ½ its old value and ¼ of the blue's old value. This produces a green slightly tinged with blue. There is variation in the amount of blue because of discrepancies among the distributions. The new blue RF takes ½ its old value. This produces a perfect blue but with half the sample size. While the new blue RF is show as blue in its entirety, its maximum extent is internally unchanged, so while the RF's GUI makes it appear as if extrapolation occurs, this is not the case if the Emergic Value is examined – it is exactly the same as before but with half the sample size. We have simply not shown the distributions for the X and Y coordinates in this example. Indeed, while the red RF appears as a perfect lilac, the extent for the R and B colour sensitivities will differ and retain their old spatial extents.

What this example shows is that the statistics for the blue RF have been *split* across the red and green RFs, while the statistics of the grey RF has been *merged* from those of red and green's. This can cause the image to dissipate over time as shown by the change of effective colour. If the eye moves back, the statistics that have no variance will preserve information, but those with variance will lead to information loss over time.

G.3   Emergic Network Architecture of ECM

The complete Emergic Network architecture of ECM – the unified cognitive model of visual filling-in – is shown in Illustration 278. Its embodied configuration depends on the

development parameters, but what it computes – it buFunctionality – remains invariant

and parameter free. Its epistemic behaviour at the macro level is dependent both on its

configuration, and on the interaction with the environment.
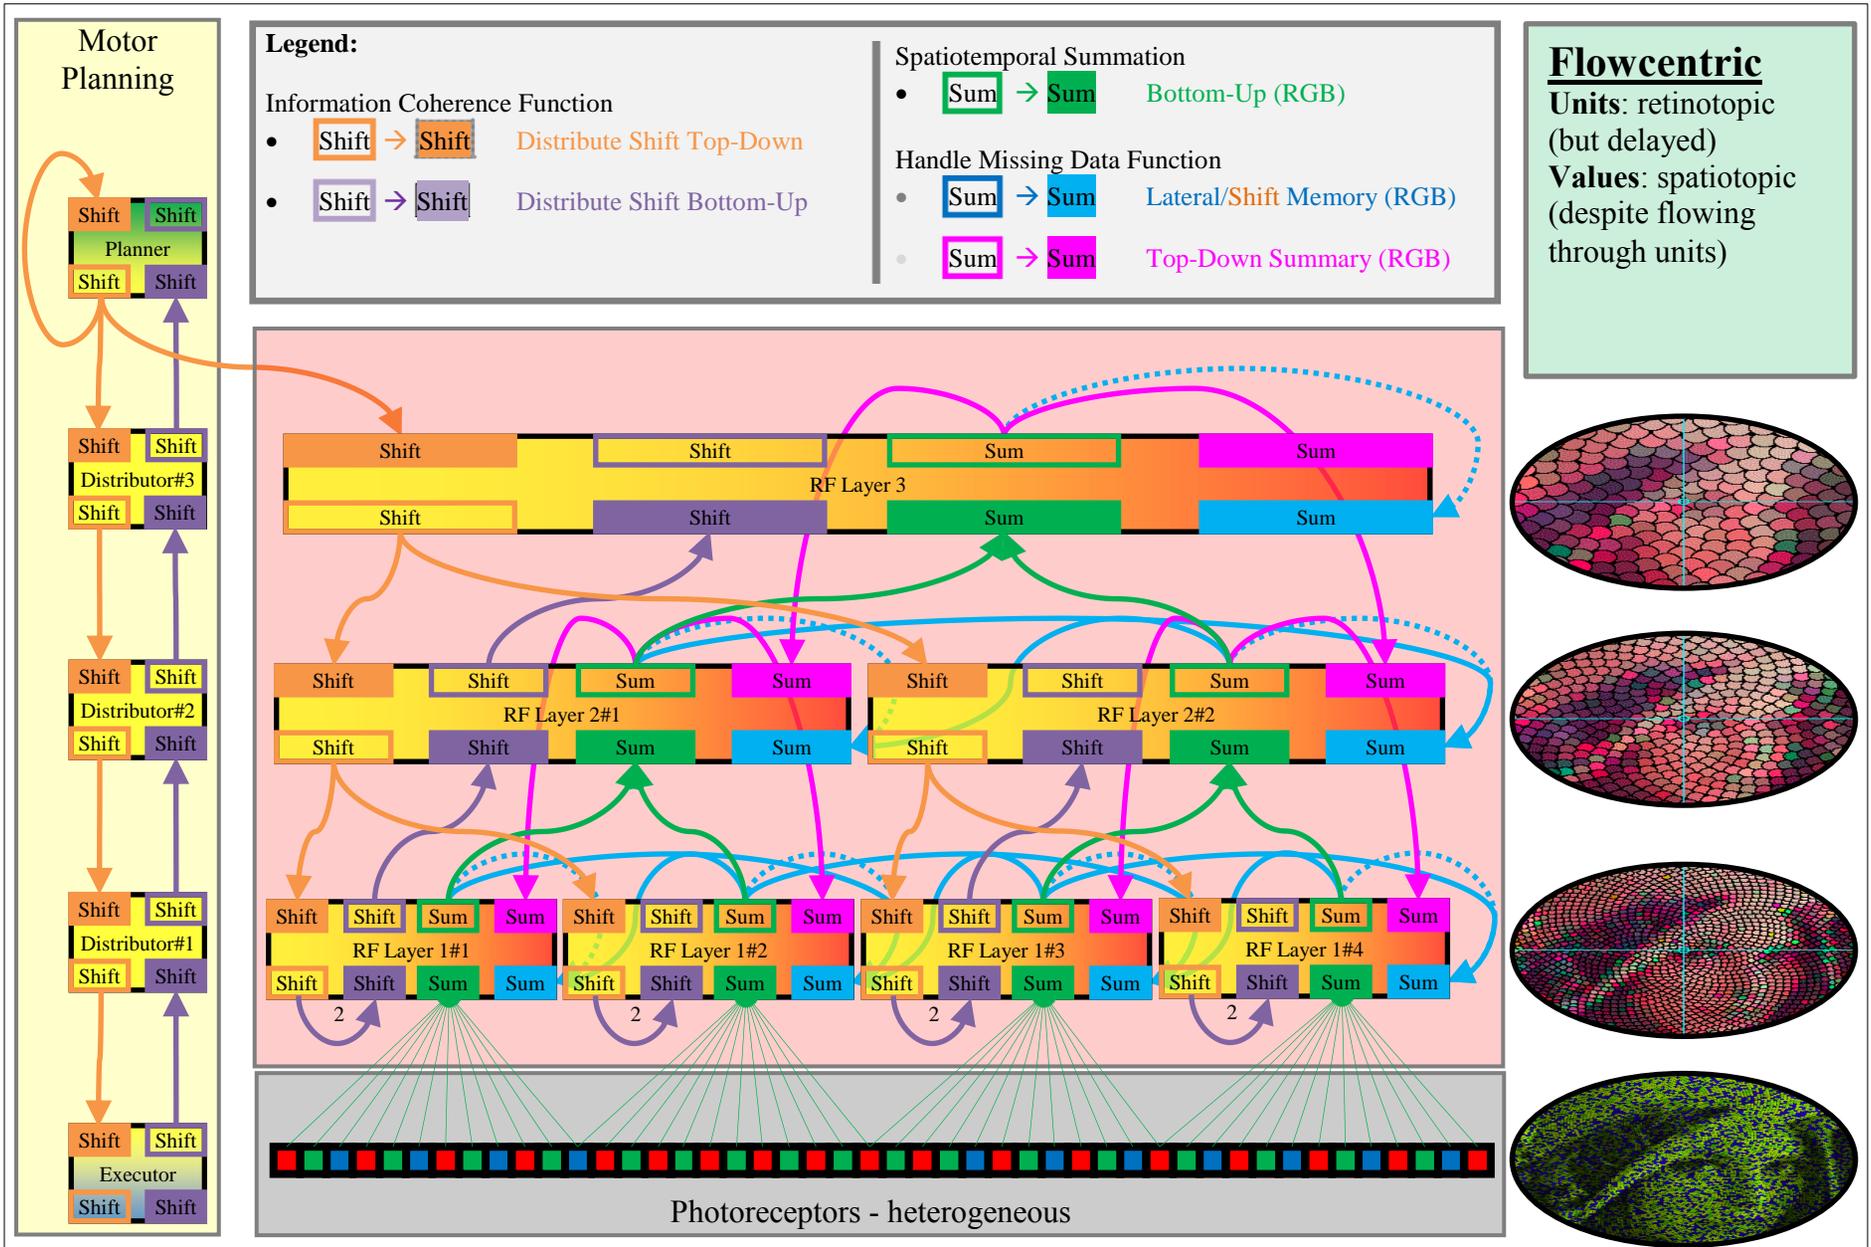
Illustration 278: Emergic Network architecture of ECM – A unified cognitive model of visual filling-in

360

G.4 Neurophysiological Support

In this section we discuss the neurophysiological support for our architecture.

G.4.1 Distributing shift information

While ECM distributes shift information in advance through the RF hierarchy, this need not represent synaptic connections. For example, the shift information could be distributed by propagating/traveling waves (Ermentrout & Kleinfeld, 2001; Sato, Nauhaus, & Carandini, 2012; Wu, Xiaoying Huang, & Chuan Zhang, 2008) and be detected by electrical gap junctions (Han & Massey, 2005). The putative mechanism is via an efference copy (Bridgeman, 2007; Imamizu, 2010; Keith et al., 2010; Perrone & Krauzlis, 2008) of a corollary discharge (Colby, Berman, Heiser, & Saunders, 2005; Crapse & Sommer, 2008; Poulet & Hedwig, 2007). There are many kinds of eye position signals and distribution pathways (Morris, Kubischik, Hoffmann, Krekelberg, & Bremmer, 2012). For example, the central thalamus distributes signals decomposed into horizontal and vertical components to the cortex (Tanaka, 2007). In distribution, there is nothing novel in ECM.

G.4.2 Shifting behaviour

That information is effectively shifted and sometimes in advance has extensive neurobiological support in terms of *behaviour* (Burr, Cicchini, Binda, & Morrone, 2011; Burr & Morrone, 2011; Duhamel et al., 1992; Hall & Colby, 2011; Hamker & Ziesche, 2011; Merriam et al., 2007; N. A. Parks & Corballis, 2010; Schneegans & Schöner, 2012). Nevertheless two mechanistic interpretations are possible. Either the information moved, or the receptive field moved. The current name for the behaviour is neural remapping, implying that there is something about the receptive field that has changed,

while visual content continues in a simple feed forward path. That the mechanism is based on moving spatially tagged information is unique to ECM.

G.4.3 Tagging surface colour information by spatial extent

Much of visual processing is considered as retinotopic. That means that the location of a neuron in the brain indexes its location relative to the retina. The location of the neuron is its "place code" or "labeled line" code – a *neurocentric* interpretation. Our model is novel in multiplexing spatiotopic information over the visual signals in an infocentric or *flowcentric* interpretation. Is this even possible?

The standard visual signal, e.g., amount of surface redness, is considered as rate encoded. However, neural spikes have many other characteristics such as phase. It is possible to nest codes within spike trains. For auditory information, it was found that the temporal spike-phase code compliments the spatial label (Kayser et al., 2009).

The whiskers of mice are arranged in a spatiotopic manner. Nevertheless the location of contacted objects is encoded via three orthogonal mechanisms. Vertical coordinates by the line code, horizontal coordinates by activation timing, and radial coordinates by activation intensity (Knutsen & Ahissar, 2009).

Throughout the visual system there are numerous neurons whose signals are modulated by eye position and gaze (Morris et al., 2012). Where the strength of modulation varies systematically (Andersen & Mountcastle, 1983) or idiosyncratically (Chang & Snyder, 2010), these are called gain fields (Blohm & Crawford, 2009). They are ubiquitous.

While not coding position explicitly, they are considered as representing space via distributed coding (Andersen, 1993).

G.4.4 Dendritic processing

Ours is not the first model to require dendritic processing. See for example, (Quaia, Optican, & Goldberg, 1998). However, our functional requirement is unique. A *distribution* is broadcast on a local network so that its portions can be summed up into new distributions consistent with the spatiotopic extent of the RFs and the original distribution. RFs in the "wrong" location will simply ignore the information.

The existence of nested and orthogonal coding implies the need for dendritic processing. However, the field of dendritic processing is extremely young (Cuntz, 2012; Destexhe, 2010; Gulledge, Carnevale, & Stuart, 2012; London & Häusser, 2005; Yuste, 2011), and applying such research to our needs and furthering such research with our theoretical requirements is for future research. We can speculate.

As far as RFs selecting one set of signals, and ignoring other, (Morita, Okada, & Aihara, 2007) have a winners take all model of selection based on dendritic nonlinearity. If propagating waves can signal spatiotopic coordinates, they can be used to form an AND gate on other cortical inputs (Ang, Carlson, & Coulter, 2005). As far as summing up a new distribution, every local broadcaster has a potentially different spatiotemporal pattern of dendritic inhibition that can drastically affect the output of a neuron (L. Palmer, Murayama, & Larkum, 2012). Excitatory integration involves active dendritic processes (Magee, 2000) that could be harnessed.

While a single artificial neuron with recurrence is Turing Complete it is often simplified to represent a single computation. Not much can be done with only a few such neurons at a given cortical site. However, the computational power of a unit grows with the number of nonlinear subunits (Polsky, Mel, & Schiller, 2004), so the possibilities inherent within dendritic processing can be immense. It remains to be seen whether our system level application finds biological plausibility.

Appendix H: Emergic Simulation System

The Emergic Simulation System (ESS) runs experimental paradigms on human participants and cognitive simulations alike. We sometimes call them *tests* when run on simulated *persons*. An experimental paradigm is a carefully controlled and measured two way real-time interaction between a *person* and their environment. Typically each

Illustration 279: Emergic Simulation System (ESS) Context

paradigm has a unique *World* controller for creating specific stimuli and manipulating the environment (see Dombeck & Reiser, 2012). Each world can be thought of as a unique virtual reality video game presented on a computer screen.

Currently, the only interactions supported for human participants are for them to move their gaze across the computer monitor and to blink their eyes. These might not be thought of as interactive, but they do change the stimulus that falls on the retina – the participant is changing the part of the environment they are engaging with, so they are indeed considered as

Illustration 280: *Person*/environment interaction via relative eye motion

interacting with their environment. Such interactions are more obvious when the visual stimulus contains objects in motion and a *person* needs to track such objects as in Illustration 280, or maintain stable percepts when their eyes move over a static scene. Such *interactions* are not representational.

Of course, the intent is for ESS to run equivalent tests on simulated *persons* so that the computational model can be validated against human data. To this end, an Emergic Cognitive Model (ECM) is "grown" and embodied within a *Person* controller as shown in Illustration 279 and Illustration 281, with each eye connected to a different *View* of the world. The views are situated, can shift gaze across the computer monitor, and blink as required. Visual content is dynamically mapped onto the views to retain non-representational interaction equivalence. While ESS supports stereoscopic vision, currently, ECM only models monocular or cyclopean vision.

The one wrinkle in the entire system is that while video games on computers can execute in real-time, visual cognitive models are extremely CPU intensive. Therefore, to ensure that the non-representational interaction between *Person* and *World* remains accurate, both are run in simulated time. This requires a *time synchronization* mechanism to distribute *ticks* of time appropriately between *Worlds*, *Views* and *Persons*. Our solution to all of these components will be described in this chapter in a *reference manual* style introduced by their respective GUI representations.

In summary, we prefer to think of each experimental paradigm or stimuli configuration as having a *person* immersed in a different *world* and playing a different computer game, with each eye receiving a slightly different *world view*. See Illustration 281.
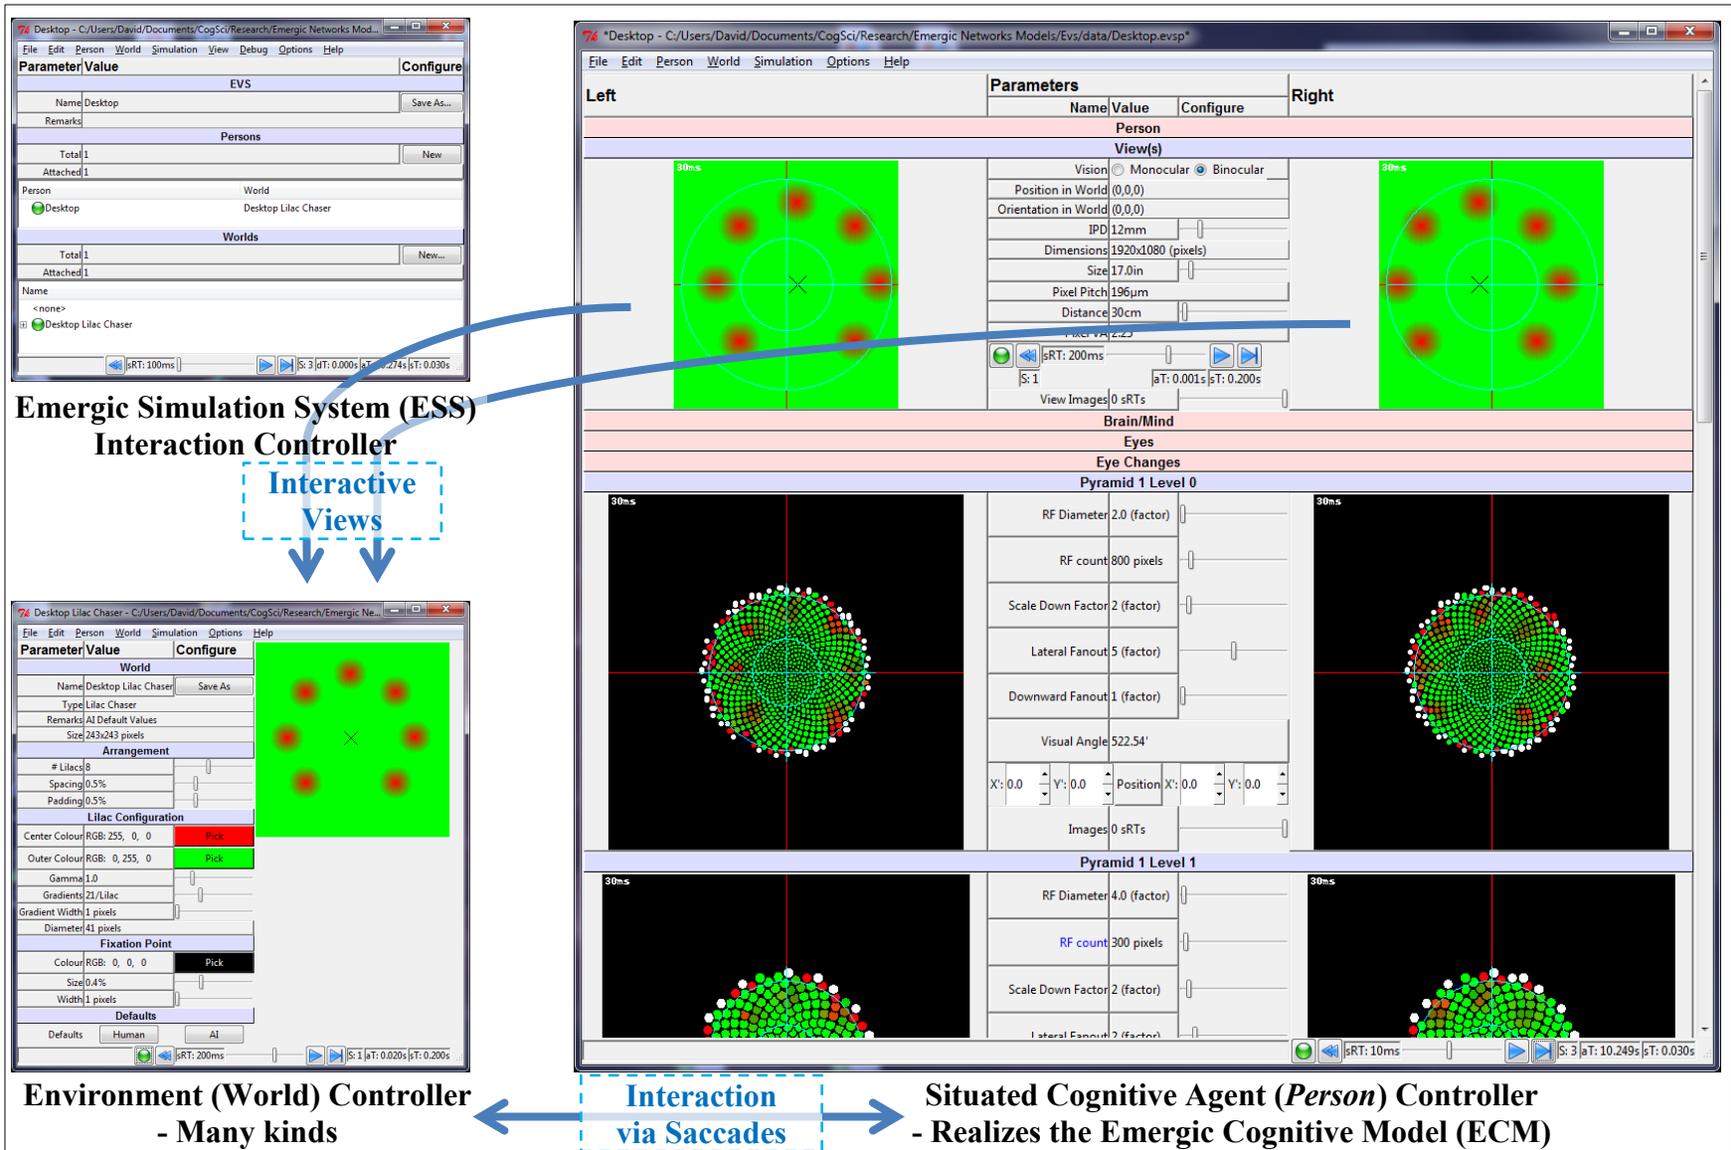
Illustration 281: ESS, World, *Person* and View interactions

Emergic Simulation System (ESS)
Interaction Controller

Interactive Views

Environment (World) Controller
- Many kinds

Interaction via Saccades

Situated Cognitive Agent (*Person*) Controller
- Realizes the Emergic Cognitive Model (ECM)

H.1 Time Synchronization

ESS integrates Worlds, Views and *Persons* in a non-representational manner, but each could be realized via different technologies running at different time frames. Our time synchronization solution is to adopt a client/server architecture whereby a server distributes ticks to all clients based on their rate requirements. In this section we describe the time synchronization solution and its abstract GUI and use by ESS.

H.1.1 GUI based introduction

Our time synchronization solution adopts a client/server architecture whereby a server distributes ticks to all clients based on their rate requirements. The system is best described via its user interface (Illustration 282) where each element is identified by a number and described next.
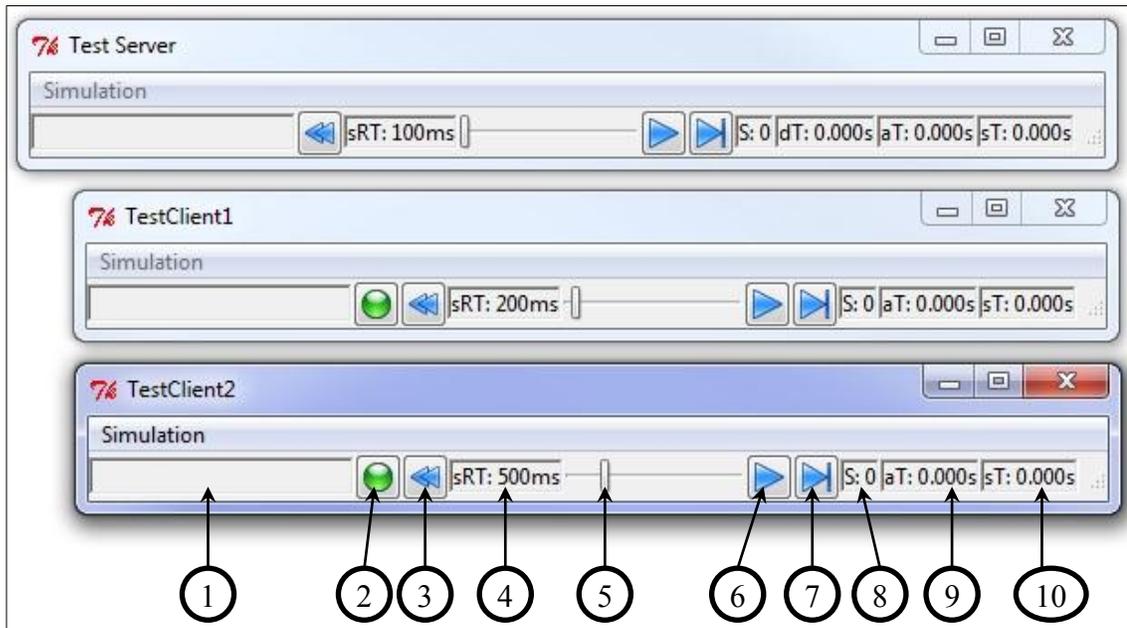
Illustration 282: Client-server time distribution GUI

H.1.1.1　Simulated reaction time value (4)

Displays the simulated reaction time. For the server, this sets the base time/tick rate, for clients, their time/tick rate are ideally integral multiples of the base rate.

H.1.1.2　Simulated reaction time control (5)

Sliding this control changes the simulation reaction time value.

H.1.2　Use by ESS

A summary of how the Emergic Simulation Systems makes use of this synchronization mechanism is described next. ESS is a time distribution server that contains three classes of time distribution clients. A *world* represents a stimuli generating subsystem such as a game server, or a lilac chaser controller
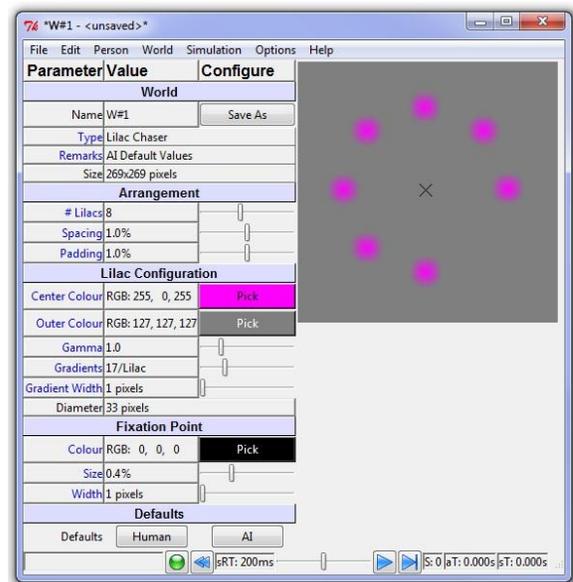
Illustration 283: Lilac chaser world controller

(Illustration 283). They are technology dependent. Multiple worlds can run at the same time. A *person* represent the brain/mind of a virtual cognitive agent such as the Emergic Cognitive Model based on the Emergic Network. Multiple *persons* can inhabit each world. A *view* represents a single two dimensional view over a world. It performs the translation service from a technology dependent world environment, to generic/internal pixels that can be mapped onto photoreceptors. A *person* would have two views, one per eyeball. The Emergic Simulation System simply creates prototype worlds, generates brains for people out of a set of parameters, and connects their views to the appropriate worlds. The Emergic Simulation System maintains a list of these clients such that if they

all need to be ticked simultaneously, they are ticked in the following order: worlds, views then *persons*. When a *person* is ticked, the misaligned internal pixel format is appropriately mapped to values sent to each photoreceptor (for each eye), and then their entire Emergic Network brain/mind is ticked.

H.2   ESS controller

ESS can handle multiple *persons* with multiple *views* onto multiple *worlds*. However, all the experiments in this thesis make use of a single cyclopean *person* interacting with a single *world view* at a time. A *person* can be disengaged from a view to simulate eye-blinks or imagination, although blink parameters were recently added to *person*.
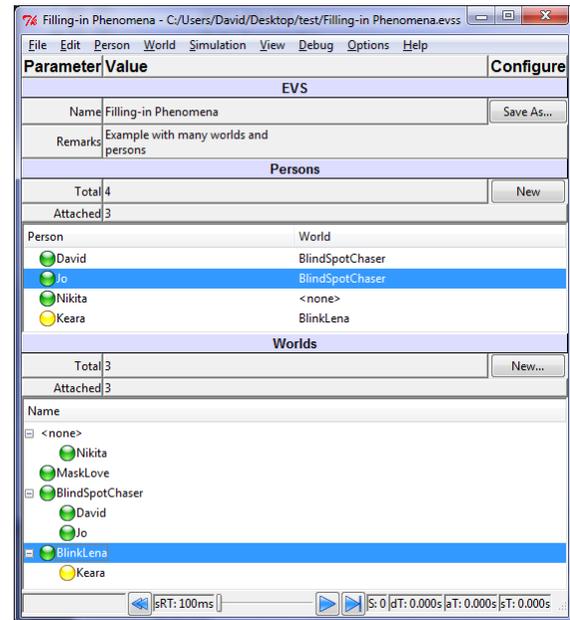
Illustration 284: ESS controller GUI

ESS is best described via its control interface (Illustration 284) where each element is described in top-down, left-to-right order. However, in this thesis we only describe the parameters needed to replicate all results and the ESS has none.

H.3   World controllers

While both *world* and *person* can be controlled or instructed symbolically, they interact in a non-representational manner, i.e., via view pixels. The world generates a dynamic view and the *person* is free to saccade across it. Many of the dynamic views at this stage of ESS development are quite static. However as ECM is continuously moving its eyes, that's ok.

H.3.1   World picker

The world picker (Illustration 285) is simply a front end for preconfigured worlds, at least one of each type.

Each type of world is described next.
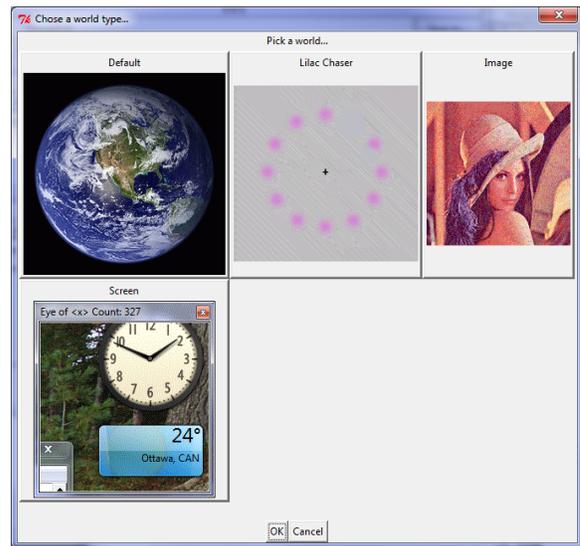
Illustration 285: World picker

372

H.3.2 Lilac chaser World

The lilac chaser is host to a wide variety of visual phenomena (Atvars & Fomins, 2009; Blue & Szirko, 2005; Winkler, McDermott, Caplovitz, & Webster, 2011). It has been placed under control (Illustration 286) for testing the Emergic Cognitive Model against a dynamic stimulus. In this case, the

Illustration 286: Lilac chaser world controller GUI

lilac chaser gap will travel one lilac every 200ms based on the sRT value.

H.3.3   Image world

The image world (Illustration 287) controls the display of a single static image. Although
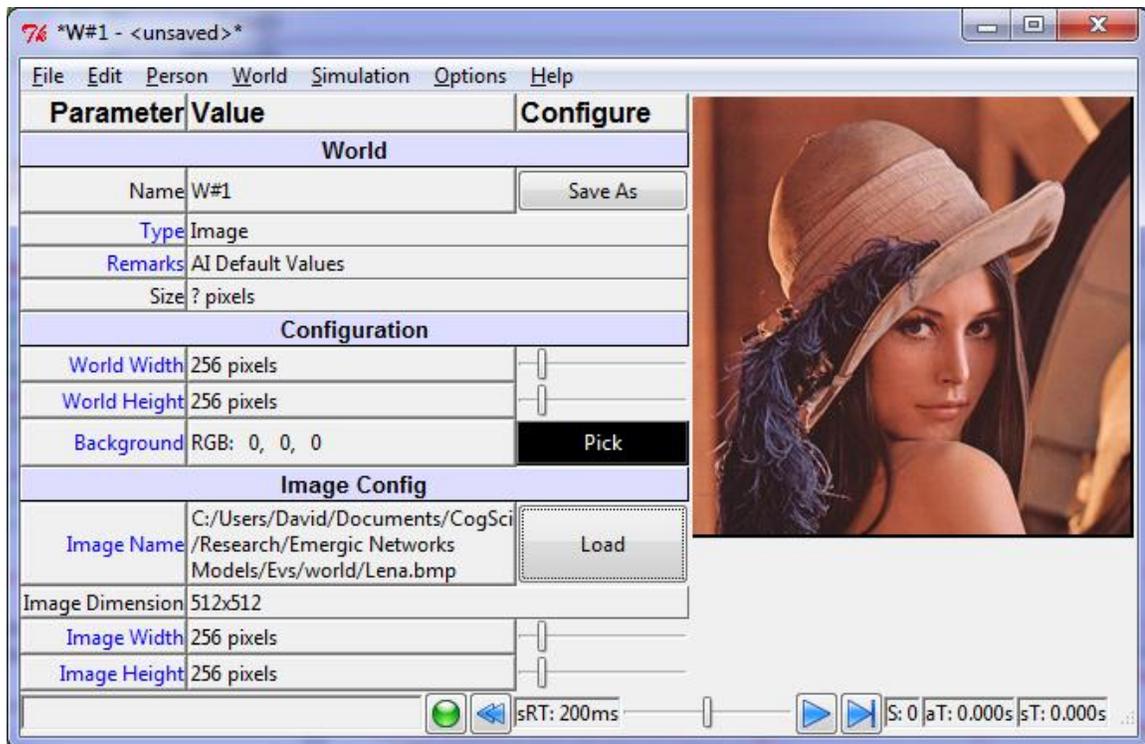
Illustration 287: Image world controller GUI

the image is static, the Emergic Cognitive Model still moves its eyes (and view) across

the image. The sRT is not relevant for static images.

H.3.4 Screen world

The screen world (Illustration 288) can capture any image on the desktop and it can move
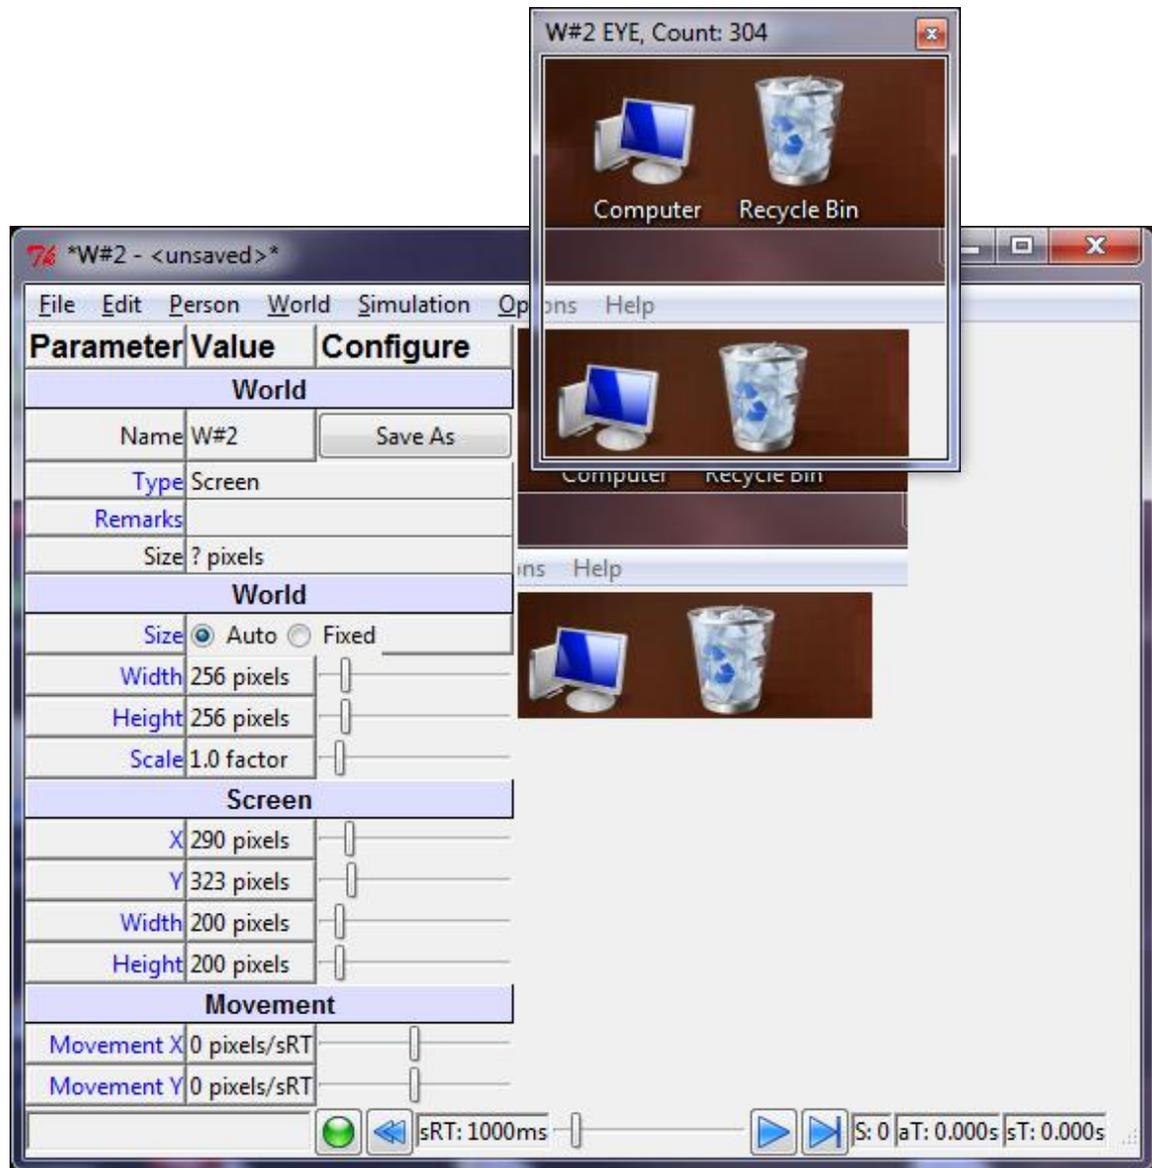
Illustration 288: Screen world controller GUIs

its screen capture camera or EYE in a regular fashion. In this example, the EYE is

recurrently placed over its own image. The images being captured should be static

because although the screen world can be synchronized to run in simulated time, the

images being captured cannot be slowed down. Typically, a static image is created in any

paint program of choice, and then the camera is position above the picture. For example, an arrow could be moved through the blind spot via this world.

H.4 *Person* controller

The Emergic Cognitive Model (ECM) is grown (e.g., Koene et al., 2009) and embodied (see M. L. Anderson, 2003; Kiverstein & Clark, 2009; Sloman, 2009) within a *person* controller. It is then situated so that it can move its eyes across different views of a world, and it can blink. This completes the perception-action link necessary for interacting with the environment (Bridgeman & Tseng, 2011). In this section, we describe the *person* controller via its GUI (Illustration 289), but the ECM and its development program are described on page 323.
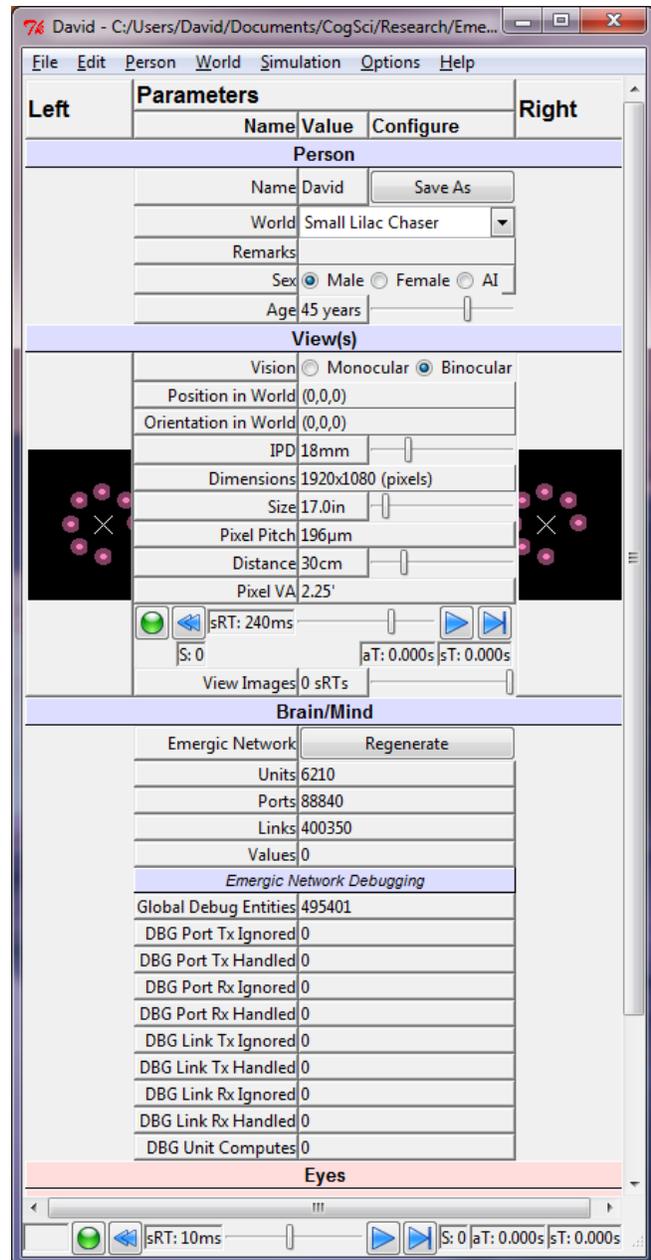
Illustration 289: *Person* controller GUI

Each element of the GUI will be described in a top-down, left-to-right order.

H.4.1 Parameter Name/Value/Configure; Left/Right

A set of parameter name/value pairs. If an item is followed by a (name) in parenthesis, this indicates the paramter

Illustration 290: *Person* **parameters**

name used when storing the value in the configuration file. If a parameter value can be changed, one can often click on the value to edit, or use a control in the *Configure* column. The parameter names turn blue if they have been changed and not saved; they turn red if in error.

The content of these parameter files are shown in the supplement for all experiments (or tests) in this thesis to precisely specify the paradigm or protocol.

The Left/Right columns show information about both brain hemispheres, eyes and views of ECM for a stereoscopic *person*. All the tests in this thesis are for monocular/cyclopean *persons* and subsequent GUIs will only show the Left column.

H.4.1.1 *Person* section

A set of descriptive parameters about this embodied and situated *person* instance of ECM. This section can be collapsed.
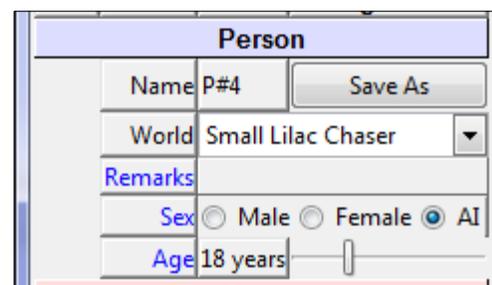
Illustration 291: *Person Person* **parameters**

H.4.1.1.1 Name

Name is a static descriptive parameter identifying this *person* from others (based on the configuration file name). A given experimental stimuli (or world) can be viewed by

378

multiple *persons* at the same time. The 'Save As' button save the current configuration into a new file and changes the value of the name parameter.

H.4.1.1.2   World

World is dynamic parameter identifying which world this *person* is currently interacting with. If set, the left and or right views will be updated. If not set, the views would be dark. Imagination can be simulated by setting the world to <none>. The list contains only currently open worlds, but the first three values are <none>, <new> and <open>.

Use <none> to remove the view from any world. Use <new> to create a new world to be viewed. Use <open> to open an existing world to be viewed.

H.4.1.1.3   Remarks (note)

Remarks is a static descriptive parameter with textual information to be filled in by the researcher. It typically explicates the aspects of cognition being studied.

H.4.1.1.4   Sex (sex)

Sex is a static descriptive parameter. ECM makes no distinction based on sex but the researcher is free to create categories of parameters related to sex.

H.4.1.1.5   Age (age)

Age is a static descriptive parameter. ECM makes no distinctions based on age but the researcher is free to create categories of parameters related to age. Note that the simulation time parameter (sT) identifies how long the *person* has lived in the current

simulation. While learning is incorporated in the dynamics, parameters do not change with age.

H.4.1.2 View(s) section

Views is a section containing parameters relevant to the visuospatial modality. It is shown here merely to provide a visual look to the bilateral organization of *person*. Note that the
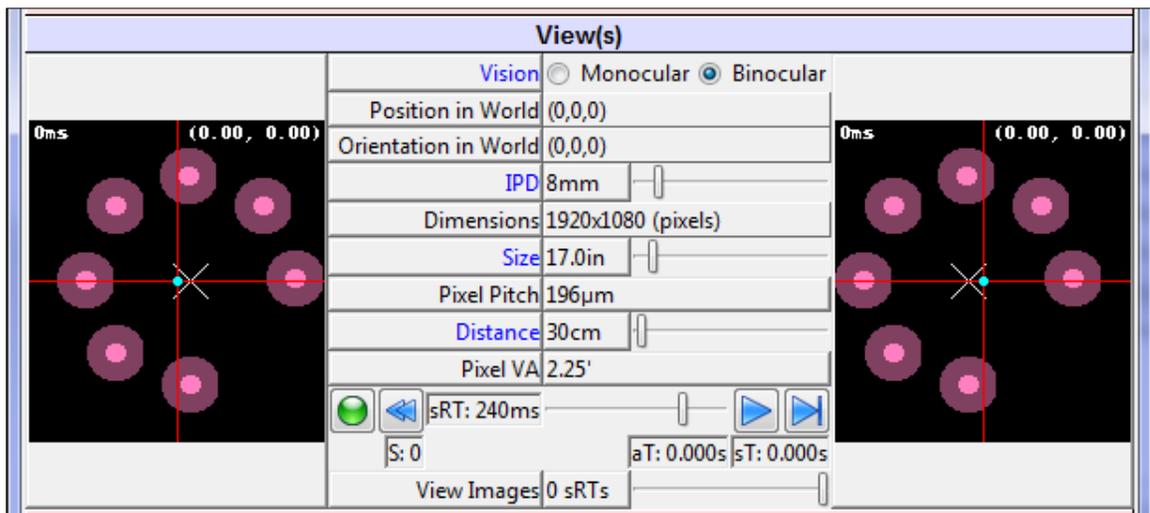
Illustration 292: *Person* **View(s) parameters**

left and right eyes have slightly different views of a world due to the Inter-Pupillary Distance (IPD) parameter (Dodgson, 2004). While ECM creates cognitive structures in two hemifields, they are not joined as that is not the focus of our research. Typically, under monocular/cyclopean vision (González, Weinstock, & Steinbach, 2007; Schreuder, 2008), only the left hemifield is generated and the *Right* column is collapsed.

H.4.1.2.1 Left/Right Images

The left and right view of the world are
separated by the IPD parameter which is
usually zero for monocular visual modelling.
The last 32 frames are maintained. Right-click
on the image to save the current frame as a GIF
file, or all frames as an animated GIF file.
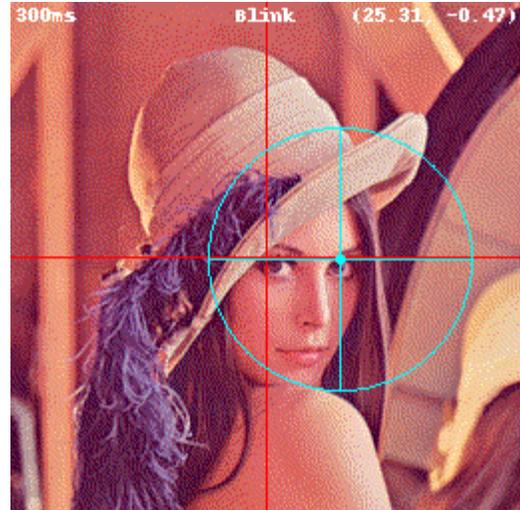
Illustration 293: View with overlay

Although the view of the world is based on the view's sRT, an overlay is added onto it as

shown in Illustration 293, based on the *person's* sRT. This allows demonstrating the

model as a set of synchronized animated GIFs. The overlay information is as follows.

The overlay's sRT is shown in the top left.

Blink is shown in the middle top only during parameterized blink periods, but not when

the *person* is not viewing a world.

The eye's position over the view is shown in the top right, as well as by the cyan

coloured circle. There is both an inner circle (small in this case) and an outer circle

showing the foveola and fovea regions.

The cyan coloured axes and retina represent the retinotopic coordinates.

The red coloured axes represent the spatiotopic coordinates, and are unaffected by eye motion.

If the mouse is hovered over the view, its (x,y) coordinates will be shown in the status bar. A single click will make that a target for the next saccade; a double click will actually move the eye.

H.4.1.2.2   Vision (eyeCount)

Vision is either monocular or binocular.

H.4.1.2.3   Position in World (wPos)

This parameter informs on the location of the view within the world.

H.4.1.2.4   Orientation in World (wOrient)

This parameter informs on the orientation of the view within the world.

H.4.1.2.5   IPD (ipd)

the Inter-Pupillary Distance (IPD) parameter (Dodgson, 2004) specifies how far apart the two eyes are. It is zero for monocular models.

H.4.1.2.6   Dimensions

The *dimensions* of the actual computer monitor are shown.

H.4.1.2.7   Size (vSize)

The diagonal screen size of the actual computer monitor must be set.

H.4.1.2.8 Pixel Pitch (ppitch)

The pixel pitch is determined by the previous parameters.

H.4.1.2.9 Distance (viewEyeZ)

The *distance* parameter specified the distance between the eye and the view (virtual computer screen). It allows one to zoom into the picture.

H.4.1.2.10 Pixel VA (pixVa)

This parameter shows the calculated visual angle of a single pixel.

H.4.1.2.11 Attach/Detach

The view can be detached and re-attached from time synchronization.

H.4.1.2.12 Rewind

Unused.

H.4.1.2.13 sRT value

The view's sRT value is typically the same as the world's.

H.4.1.2.14 sRT control

This control allows the sRT value to be changed.

H.4.1.2.15 Play/Pause

Continue or pause the experiment. If the view is attached, all other attached items will be played or paused appropriately.

H.4.1.2.16   S

The number of sRT tick steps this view has undergone.

H.4.1.2.17   aT

The amount of actual time attributed to processing this view.

H.4.1.2.18   sT

The amount of simulated time attributed to processing this view,

H.4.1.2.19   View Images

A control to choose which of the last 32 view images to display.

H.4.1.3   Brain/Mind section

This section contains a set of dynamic
parameters describing the Emergic Network
making up the brain/mind of this *person*. This
section can be collapsed.

H.4.1.3.1   Emergic Network

The regenerate button will generate a new
brain/mind for this *person* based on the
development parameters. Changes to some of
these parameters require regeneration. Even

Brain/Mind	
Emergic Network	Regenerate
Units	3105
Ports	44420
Links	200175
Values	0
Emergic Network Debugging	
Global Debug Entities	247701
DBG Port Tx Ignored	0
DBG Port Tx Handled	0
DBG Port Rx Ignored	0
DBG Port Rx Handled	0
DBG Link Tx Ignored	0
DBG Link Tx Handled	0
DBG Link Rx Ignored	0
DBG Link Rx Handled	0
DBG Unit Computes	0

Illustration 294: *Person* **Brain/Mind parameters**

when the parameters are unchanged, each brain will differ due to stochastic components,

such as the determination of which L, M or S colour sensitivity types a photoreceptor will

be assigned.

H.4.1.3.2   Units

This is a dynamic parameter indicating how many Emergic Units are currently in the Emergic Network. Although dynamic, it does not change during the experiment. Only regeneration will affect this parameter.

H.4.1.3.3   Ports

This is a dynamic parameter indicating how many Emergic Ports are currently in the Emergic Network. Although dynamic, it does not change during the experiment. Only regeneration will affect this parameter.

H.4.1.3.4   Links

This is a dynamic parameter indicating how many Emergic Links are currently in the Emergic Network. Although dynamic, it does not change during the experiment. Only regeneration will affect this parameter.

H.4.1.3.5   Values

This is a dynamic parameter indicating how many Emergic Values are currently flowing in the Emergic Network. This parameter will change during the experiment. It is typically zero at the beginning. The first set of values can come from two sources. Firstly there is bottom-up values coming from the photoreceptors, but nothing will happen if the *person* is not viewing a world. Secondly, there is planned eye motion that is distributed top-down. In special cases, the software can be engineered to inject values, e.g., to test imagination (p299).

H.4.1.3.6   Emergic Network Debugging subsection

This is a set of dynamic parameters useful for debugging the Emergic Network. They are not described and can change without notice.

H.4.1.4   Eyes section

This section contains the static development parameters and dynamic informative and control information for each eye. The section can be collapsed.
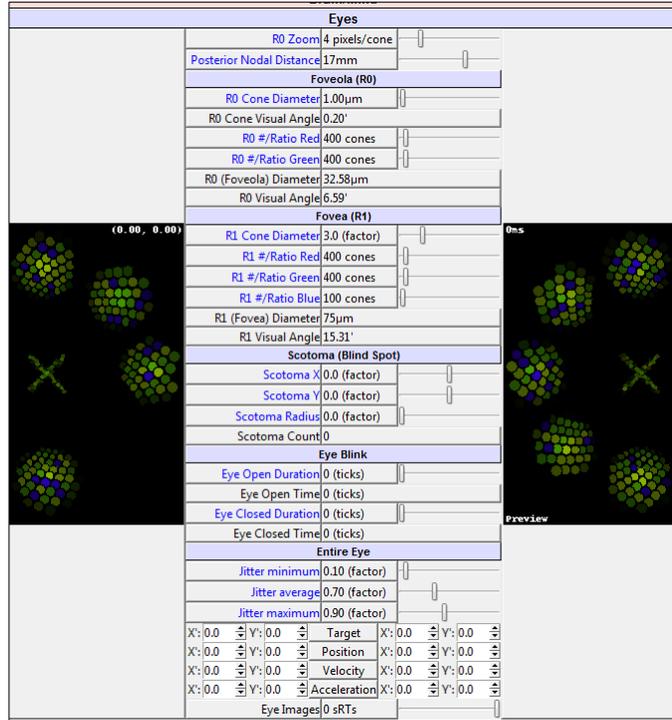
Illustration 295: *Person* **Eyes parameters**

H.4.1.4.1   Left/Right Image

These images show the Photoreceptor emergic units. They can be shown in several modes. If the *person* is not connected to any view than all the photoreceptors will be previewed fully lit in their false RGB colours (Illustration 296). When viewing a world, the photoreceptors will show the appropriate LMS colour space activations (Illustration 297).

Ordinarily, the photoreceptors show what they last received from the view. However,
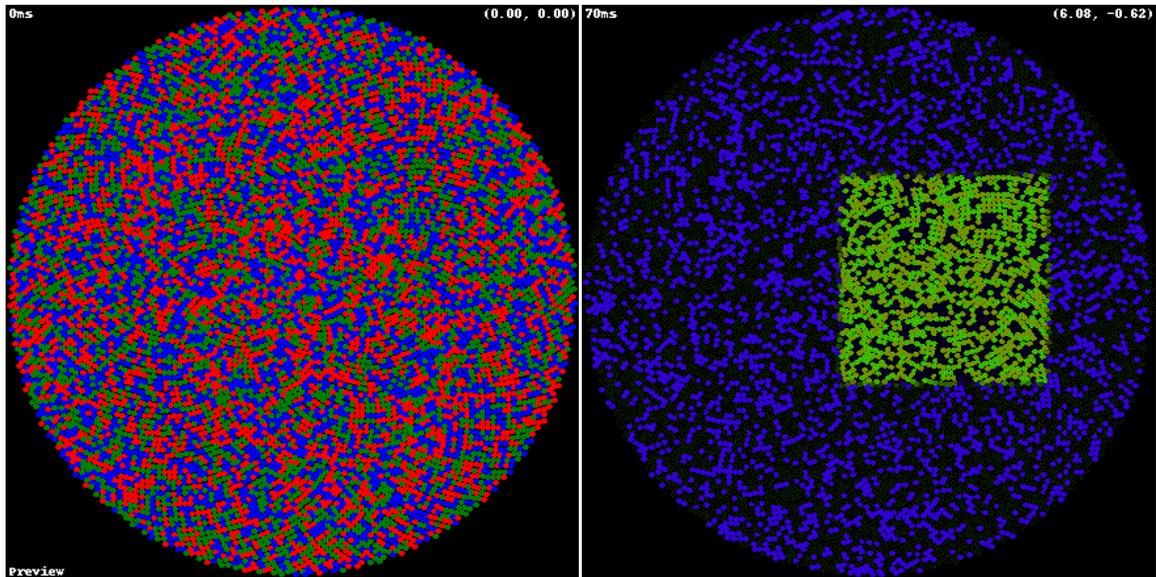
Illustration 296: Photoreceptors viewing no world showing false RGB colours

Illustration 297: Photoreceptors viewing a world showing LMS colour activations

when the simulation is paused, they can *preview* the view if the eye is moved around.

Their overlays include the current sRT (top-left), x-y coordinates (top-right), 'Preview' (bottom-left) if in preview mode, and 'Blink' (centered) if the eye is currently shut and all photoreceptors will appear as black.

Right clicking on an image allows saving the current one as a GIF or the last 32 as an animated GIF.

H.4.1.4.2   R0 Zoom (r0ConePix)

This parameter controls the GUI zoom and specifies how many pixel make up the diameter of the smallest R0 region cone. A value of 6 or greater allows the cone to be displayed with a black border on the GUI.

H.4.1.4.3   Posterior Nodal Distance (nodalPt)

This static parameter specifies the distance between the retina and cornea lens midpoint, .i.e., the focal length.

H.4.1.4.4   Foveola (R0) subsection

The foveola is the innermost region (R0) of the retina. This section contains related parameters. The section can be collapsed.

Foveola (R0)	
R0 Cone Diameter	1.00µm
R0 Cone Visual Angle	0.20'
R0 #/Ratio Red	400 cones
R0 #/Ratio Green	400 cones
R0 (Foveola) Diameter	32.58µm
R0 Visual Angle	6.59'

Illustration 298: *Person* **Foveola (R0) parameters**

H.4.1.4.5   R0 Cone Diameter (r0ConeD)

This static development parameter specifies the diameter of the smallest cone in the foveola (region 0). It is 0.5um for humans.

H.4.1.4.6   R0 Cone Visual Angle (r0ConeVa)

This dynamic parameter indicates the developed visual angle of the smallest cone in the foveola (region 0).

H.4.1.4.7   R0 #/Ratio Red (r0Red)

This static development parameter indicates the number/ratio of L (red) cones to stochastically allocate within the foveola (region 0). There are no S (blue) cones in this region.

H.4.1.4.8   R0 #/Ratio Green (r0Green)

This static development parameter indicates the number/ratio of M (green) cones to stochastically allocate within the foveola (region 0). There are no S (blue) cones in this region.

H.4.1.4.9   R0 (Foveola) Diameter

This dynamic parameter indicates the developed diameter of the foveola (region 0). It is typically 200um for humans.

H.4.1.4.10   R0 Visual Angle

This dynamic parameter indicates the developed visual angle of the foveola (region 0).

H.4.1.4.11   Fovea (R1) subsection

The fovea is the innermost region (R1) of the retina. This section contains related parameters. The section can be collapsed.

Fovea (R1)		
R1 Cone Diameter	3.0 (factor)	
R1 #/Ratio Red	400 cones	
R1 #/Ratio Green	400 cones	
R1 #/Ratio Blue	100 cones	
R1 (Fovea) Diameter	75µm	
R1 Visual Angle	15.31'	

Illustration 299: *Person* **Fovea (R1) parameters**

H.4.1.4.12   R1 Cone Diameter (r1ConeD)

This static development parameter indicates the size of the outermost cone within the fovea region (R1) as a factor of the smallest cones within the foveola region (R0).

H.4.1.4.13   R1 #/Ratio Red (r1Red)

This static development parameter indicates the number/ratio of L (red) cones to stochastically allocate within the fovea (region 1).

H.4.1.4.14   R1 #/Ratio Green (r1Green)

This static development parameter indicates the number/ratio of M (green) cones to stochastically allocate within the fovea (region 1).

H.4.1.4.15   R1 #/Ratio Blue (r1Blue)

This static development parameter indicates the number/ratio of S (blue) cones to stochastically allocate within the fovea (region 1).

H.4.1.4.16   R1 (Fovea) Diameter

This dynamic parameter indicates the developed diameter of the fovea (region 1). It is 1mm for humans and uses up 50% of the visual cortex.

H.4.1.4.17   R1 Visual Angle

This dynamic parameter indicates the developed visual angle of the fovea (region 1).

H.4.1.4.18   Scotoma (Blind Spot) subsection

This section contains parameters related to the blind spot – a scotoma. The section can be collapsed. Note that the full retina is developed first and then the photoreceptors within the scotoma are disabled.

Scotoma (Blind Spot)		
Scotoma X	0.0 (factor)	
Scotoma Y	0.0 (factor)	
Scotoma Radius	0.0 (factor)	
Scotoma Count	0	

Illustration 300: *Person* **Scotoma (Blind Spot) parameters**

H.4.1.4.19   Scotoma X (blindX)

This development parameter indicates the x-coordinate location of the blind spot as a factor of the smallest cones in the foveola (region 0).

H.4.1.4.20   Scotoma Y (blindY)

This development parameter indicates the y-coordinate location of the blind spot as a factor of the smallest cones in the foveola (region 0).

H.4.1.4.21   Scotoma Radius (blindR)

This development parameter indicates the radius of the blind spot as a factor of the smallest cones in the foveola (region 0). The value of zero indicates no blind spot.

H.4.1.4.22   Scotoma Count

This dynamic parameter indicates the number of photoreceptors 'killed' to develop the blind spot.

H.4.1.4.23   Eye blink subsection

This section contains parameters that control regular eye blinking. The section can be collapsed. Note that

Illustration 301: *Person* **Eye blink parameters**

eye blinks can also be manually simulated by removing this *person* from the view, and adding them back in.

H.4.1.4.24   Eye Open Duration (blink0D)

This static parameter controls the regular developed eye open duration. The value of zero indicates that the eye will remain open and not blink.

H.4.1.4.25 Eye Open Time

This parameter indicates how much time the eyes have been open.

H.4.1.4.26 Eye Closed Duration (blink1D)

This static parameter controls the regular developed eye closed duration. The value of zero indicates that the eye will remain closed and not open again.

H.4.1.4.27 Eye Closed Time

This parameter indicates how much time the eyes have been closed.

H.4.1.4.28 Entire Eye subsection

This section contains parameters related to the overall eye. The section can be collapsed.
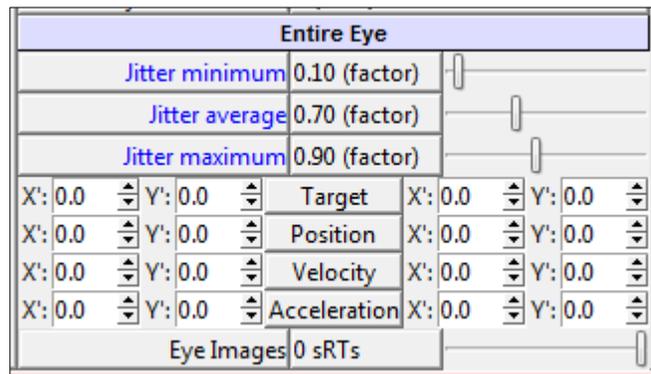
Illustration 302: *Person* **Entire Eye parameters**

H.4.1.4.29 Jitter minimum (jitMin)

This static parameter controls the developed minimum (low) jitter amount. The stochastic jitter amount depends on the triangular(low=min, high=max, mode=ave) distribution from Python.

H.4.1.4.30 Jitter average (jitAve)

This static parameter controls the developed average (mode) jitter amount. The stochastic jitter amount depends on the triangular(low=min, high=max, mode=ave) distribution from Python.

H.4.1.4.31 Jitter maximum (jitMax)

This static parameter controls the developed maximum (high) jitter amount. The stochastic jitter amount depends on the triangular(low=min, high=max, mode=ave) distribution from Python.

H.4.1.4.32 Target

These dynamic parameters control (or indicate) the x,y-coordinates of target eye motion. Each eye has its own controls. The maximum eye motion per tick based on this parameter is five of the smallest foveola cones (region 0). The target can also be set by single-clicking on the eye's view.

H.4.1.4.33 Position

These dynamic parameters indicate the current (x,y) position of each eye. This information is also available on the overlays, but note that every level shifts at its own time. The position can also be set by double-clicking on the eye's view.

H.4.1.4.34 Velocity

These dynamic parameters control (or indicate) the x,y velocity of target eye motion. Each eye has its own controls. The eye position will be increased by these amounts every tick.

H.4.1.4.35 Acceleration

These dynamic parameters control the x,y acceleration of target eye motion. Each eye has its own controls. The eye velocity will be increased by these amounts every tick.

H.4.1.4.36   Eye Images

A control to choose which of the last 32 photoreceptor images to display.

H.4.1.5   Eye Changes section

This section indicates dynamic photoreceptor changes. This section can be collapsed.
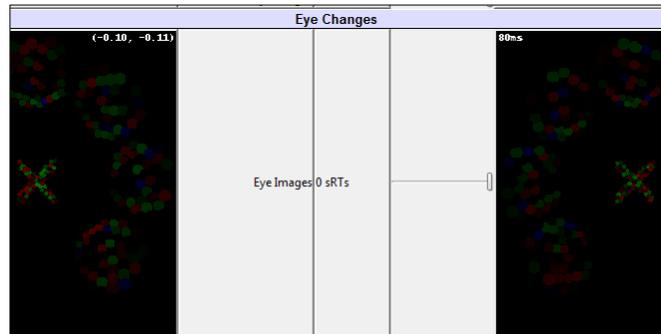
Illustration 303: *Person* **Eye Changes parameters**

H.4.1.5.1   Left/Right Image

Each photoreceptor calculates its change from its previous value. These images provide for temporal edges, but are not available until the 2nd tick. They are not yet a part of the cognitive theory within ECM. They are experimental. In particular, photoreceptors do not need to shift information, so the bigger the saccade, the larger the width of these temporal edges. We may move such change detection into the RFs where planned eye movement would generate zero edges if there was no dissipation or object motion.

H.4.1.5.2   Eye Images

A control to choose which of the last 32 photoreceptor change images to display.

H.4.1.6 Pyramid 1 Level <n> sections

This section is repeated for 3 RF levels, although they count from 0 to 2. The section can be collapsed.

H.4.1.6.1 Left/Right Image

These are similar to the

Illustration 304: *Person* **Pyramid 1 Level <n> parameters**

photoreceptor images except that they are for the RFs which have the three LMS colour values kept separate. The GUI displays them as if they were together and converted to RGB. Where all the LMS values are unknown, the RF is shown in white in which case it is ambiguous with white.

The images are twice as large due to padding. This allows the image to shift the entire RF level to keep the spatiotopic flow stable on the GUI.

Overlays include the red axes which is spatiotopic and the cyan axes which is retinotopic. In cyan is also indicated the two foveola and fovea regions of the retina. The blind spot, if any, is indicated via a black outline. The current sRT (top-left) and the levels shifted x-y coordinates (top-right) are also shown.

Right clicking on an image allows saving the current one as a GIF or the last 32 as an animated GIF.

H.4.1.6.2 RF Diameter (p1_<n>_PixD)

This static development parameter specifies the diameter of the RFs in this level as a factor of the smallest foveola cone (region 0).

H.4.1.6.3 RF count (p1_<n>_Pix)

This static development parameter specifies the number of RFs to generate within the RF level. It can cover an area larger than the photoreceptors.

H.4.1.6.4 Scale Down Factor (p1_<n>_Scale)

This static GUI parameter specifies how many times to scale down the size of the image as they are firstly doubly sized to allow for RF shifting, and secondly larger due to having areas beyond the photoreceptors.

H.4.1.6.5 Lateral Fanout (p1_<n>_FanLat)

This static development parameter specifies how much connectivity should be fanned-out to allow for shifting values flowing laterally.

H.4.1.6.6 Downward Fanout (p1_<n>_FanDwn)

This static development parameter specifies how much connectivity should be fanned-out to allow for shifting values flowing downwards. There is no upwards fan-out parameter as the system has been engineered not to require this even with shifting.

H.4.1.6.7 Visual Angle

This dynamic parameter indicates the visual angle of the developed RF level.

H.4.1.6.8   Position

This dynamic parameter indicates the amount this level has shifted in the X and Y coordinates to compensate for eye motion. This value is also indicated within the image overlay.

H.4.1.6.9   Images

A control to choose which of the last 32 RF level images to display.

Appendix I: Formulae Used

Here we list the various formulae used in the thesis or solution.

I.1   RF Diameter → RF Visual Angle

Given a receptive field diameter (d) in the retina, and an effective focal length (f), then by simple geometry, the visual angle (θ) in radians is

$$\theta = 2 \times arctan(\frac{d}{2f})$$
(1)

For example, in (Hartridge, 1922), the mean diameter of a foveal cone is estimated as d=3.2μm, while the eye's focal length is taken to be f=15mm, giving a visual angle of θ=44 seconds of arc as calculated below

$$\theta = \frac{180°}{\pi} \times \frac{60'}{1°} \times \frac{60''}{1'} \times 2 \times arctan(\frac{3.2\mu m}{2 \times 15mm} \times \frac{1mm}{1000\mu m}) = 44.003158''$$

Today, the effective focal length, also called the posterior nodal distance, is taken to be 16.67mm (Hirsch & Curcio, 1989). In the rest of the brain, e.g., in V1, this angle would need to be multiplied by the cortical magnification factor which depends on horizontal and vertical eccentricities. In this thesis, only the retinal formula is required.

I.2   RGB Colour Space → LMS Colour Space

Once a photoreceptor's sense over the RGB pixels is computed, it is converted to one of

the LMS colour space coordinates (Capilla, Malo, Luque, & Artigas, 1998; Reinhard,

Adhikhmin, Gooch, & Shirley, 2001).

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.3811 & 0.5783 & 0.0402 \\ 0.1967 & 0.7244 & 0.0782 \\ 0.0241 & 0.1288 & 0.8444 \end{bmatrix} \times \begin{bmatrix} R \\ G \\ B \end{bmatrix} \qquad (2)$$

Note that this is not corrected for device independence. Thus what a human observer

would sense from the real computer monitor and what the *Person* would sense from the

virtual monitor/*view* can be different. These quantitative differences are not our epistemic

concern.

I.3   LMS Colour Space → RGB Colour Space

The GUI converts the LMS colour space coordinates as used by the RFs to RGB colour

space coordinates as used by the computer monitor.

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 4.4687 & -3.5887 & 0.1196 \\ -1.2197 & 2.3831 & -0.1626 \\ 0.0585 & -0.2611 & 1.2057 \end{bmatrix} \times \begin{bmatrix} L \\ M \\ S \end{bmatrix} \qquad (3)$$

399

References

Abadi, R. V, Jeffery, G., & Murphy, J. S. (2011). Awareness and filling-in of the human blind spot: linking psychophysics with retinal topography. *Investigative ophthalmology & visual science*, *52*(1), 541–8. doi:10.1167/iovs.10-5910

Acker, F. (2008). New findings on unconscious versus conscious thought in decision making: additional empirical data and meta-analysis. *Judgment and Decision Making*, *3*(4), 292–303. Retrieved from http://journal.sjdm.org/vol3.4.htm

Adelson, E. H. (1983). What is iconic storage good for? *Behavioral and Brain Sciences*, *6*(01), 11. doi:10.1017/S0140525X00014291

Ahnelt, P. K. (1998). The photoreceptor mosaic. *Eye (London, England)*, *12*(Pt 3b), 531–40. doi:10.1038/eye.1998.142

Albright, T. D., Jessell, T. M., Kandel, E. R., & Posner, M. I. (2001). Progress in the neural sciences in the century after Cajal (and the mysteries that remain). *Annals of the New York Academy of Sciences*, *929*(1), 11–40. doi:10.1111/j.1749-6632.2001.tb05704.x

Alley, D., Susstrunk, S., & Herault, J. (2005). Linear demosaicing inspired by the human visual system. *IEEE Transactions on Image Processing*, *14*(4), 439–449. doi:10.1109/TIP.2004.841200

Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological science*, *16*(8), 637–43. doi:10.1111/j.1467-9280.2005.01587.x

Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track? Evidence for a resource-limited attentive tracking mechanism. *Journal of vision*, *7*(13), 14.1–10. doi:10.1167/7.13.14

Alvarez, G. A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(18), 7345–50. doi:10.1073/pnas.0808981106

Andersen, R. A. (1993). Coordinate transformations in the representation of spatial information. *Current Opinion in Neurobiology*, *3*(2), 171–176. doi:10.1016/0959-4388(93)90206-E

Andersen, R. A., & Mountcastle, V. B. (1983). The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *Journal*

of Neuroscience, *3*(3), 532–48. Retrieved from
http://www.ncbi.nlm.nih.gov/pubmed/6827308

Anderson, B. L. (2008). Transparency and Occlusion. In A. I. Basbaum, A. Kaneko, G. M. Shepherd, G. Westheimer, T. D. Albright, R. H. Masland, P. Dallos, et al. (Eds.), *The Senses: A Comprehensive Reference* (pp. 2.239–244). Elsevier. doi:10.1016/B978-012370880-9.00319-4

Anderson, C. H., & Van Essen, D. C. (1987). Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proceedings of the National Academy of Sciences of the United States of America*, *84*(17), 6297–301. Retrieved from www.ncbi.nlm.nih.gov/pubmed/3114747

Anderson, J., & Anderson, B. (1993). The Myth of Persistence of Vision Revisited. *Journal of Film and Video*, *45*(1), 3–12. Retrieved from http://web.archive.org/web/20080526105906/http://www.uca.edu/org/ccsmi/ccsmi/classicwork/Myth+Revisited.htm

Anderson, M. L. (2003). Embodied Cognition: A field guide. *Artificial Intelligence*, *149*(1), 91–130. doi:10.1016/S0004-3702(03)00054-7

Ang, C. W., Carlson, G. C., & Coulter, D. A. (2005). Hippocampal CA1 circuitry dynamically gates direct cortical inputs preferentially at theta frequencies. *Journal of Neuroscience*, *25*(42), 9567–80. doi:10.1523/JNEUROSCI.2992-05.2005

Angelucci, A., & Sainsbury, K. (2006). Contribution of feedforward thalamic afferents and corticogeniculate feedback to the spatial summation area of macaque V1 and LGN. *Journal of comparative neurology*, *498*(3), 330–51. doi:10.1002/cne.21060

Anstis, S. M. (2010). Visual filling-in. *Current Biology*, *20*(16), R664–6. doi:10.1016/j.cub.2010.06.029

Arai, I., Tanaka, M., & Tachibana, M. (2010). Active roles of electrically coupled bipolar cell network in the adult retina. *Journal of Neuroscience*, *30*(27), 9260–70. doi:10.1523/JNEUROSCI.1590-10.2010

Araragi, Y., Ito, H., & Sunaga, S. (2008). Appearance of an illusory object in the blind spot. *Perception*, *37*(8), 1301–1304. doi:10.1068/p6040

Arnold, D. H., Ong, Y., & Roseboom, W. (2009). Simple differential latencies modulate, but do not cause the flash-lag effect. *Journal of vision*, *9*(5), 4.1–8. doi:10.1167/9.5.4

Arredondo, M. A., Lebart, K., & Lane, D. (2004). Optical flow using textures. *Pattern Recognition Letters*, *25*(4), 449–457. doi:10.1016/j.patrec.2003.11.007

Arrington, K. F. (1994). The temporal dynamics of brightness filling-in. *Vision Research*, *34*(24), 3371–3387. doi:10.1016/0042-6989(94)90071-X

Assad, J. A., & Maunsell, J. H. R. (1995). Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature*, *373*(6514), 518–21. doi:10.1038/373518a0

Assisi, C., Stopfer, M., & Bazhenov, M. (2011). Using the Structure of Inhibitory Networks to Unravel Mechanisms of Spatiotemporal Patterning. *Neuron*, *69*(2), 373–386. doi:10.1016/j.neuron.2010.12.019

Atvars, U., & Fomins, S. (2009). Chaser illusion and opponent colour aftereffects. *Perception: EVCP2009: Thirty-second European Conference on Visual Perception*, *38*(ECVP Supplement), 185. Retrieved from http://www.perceptionweb.com/abstract.cgi?id=v090826

Awater, H., Kerlin, J. R., Evans, K. K., & Tong, F. (2005). Cortical representation of space around the blind spot. *Journal of neurophysiology*, *94*(5), 3314–24. doi:10.1152/jn.01330.2004

Aymoz, C., & Viviani, P. (2004). Perceptual asynchronies for biological and non-biological visual events. *Vision research*, *44*(13), 1547–63. doi:10.1016/j.visres.2003.12.011

Baccus, S. A., Olveczky, B. P., Manu, M., & Meister, M. (2008). A retinal circuit that computes object motion. *Journal of Neuroscience*, *28*(27), 6807–17. doi:10.1523/JNEUROSCI.4206-07.2008

Baijal, S. (2009). How attention affects border ownership in early visual areas. *Journal of Neuroscience*, *29*(20), 6392–3. doi:10.1523/JNEUROSCI.1156-09.2009

Bair, W., Cavanaugh, J. R., Smith, M. a, & Movshon, J. A. (2002). The timing of response onset and offset in macaque visual neurons. *Journal of Neuroscience*, *22*(8), 3189–205. doi:20026257

Balleine, B. W. (2007). The Neural Basis of Choice and Decision Making. *Journal of Neuroscience*, *27*(31), 8159–8160. doi:10.1523/JNEUROSCI.1939-07.2007

Baloch, A. A., & Grossberg, S. (1997). A neural model of high-level motion processing: Line motion and formotion dynamics. *Vision Research*, *37*(21), 3037–3059. doi:10.1016/S0042-6989(97)00103-X

Baraas, R. C., Kulikowski, J. J., & Muldoon, M. R. (2010). Bar-like S-cone stimuli reveal the importance of an intermediate temporal filter. *Journal of the Optical Society of America A*, *27*(4), 766. doi:10.1364/JOSAA.27.000766

Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, *1*(4), 371–94. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/4377168

Barnes, G. R., & Collins, C. J. S. (2008a). Internally generated smooth eye movement: its dynamic characteristics and role in randomised and predictable pursuit. *Progress in Brain Research, Vol. 171: Using Eye Movements as an Experimental Probe of Brain function - A Symposium in Honor of Jean Büttner-Ennever*, *171*, 441–9. doi:10.1016/S0079-6123(08)00665-1

Barnes, G. R., & Collins, C. J. S. (2008b). The influence of briefly presented randomized target motion on the extraretinal component of ocular pursuit. *Journal of neurophysiology*, *99*(2), 831–42. doi:10.1152/jn.01033.2007

Barrett, A. B., & Seth, A. K. (2011). Practical measures of integrated information for time-series data. *PLoS computational biology*, *7*(1), e1001052. doi:10.1371/journal.pcbi.1001052

Bartels, A., & Zeki, S. M. (2006). The temporal order of binding visual attributes. *Vision research*, *46*(14), 2280–6. doi:10.1016/j.visres.2005.11.017

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of vision*, *9*(10), 7.1–11. doi:10.1167/9.10.7

Bechtel, W., & Abrahamsen, A. (2010). Understanding the Brain as an Endogenously Active Mechanism. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 31–36. Retrieved from http://cwl.cogsci.rpi.edu/cogsci10/cogsci10_proceedings/papers/0004/index.html

Beck, D. M., & Kastner, S. (2009). Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision research*, *49*(10), 1154–65. doi:10.1016/j.visres.2008.07.012

Becker, W., & Fuchs, A. F. (1985). Prediction in the oculomotor system: smooth pursuit during transient disappearance of a visual target. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *57*(3), 562–75. doi:10.1007/BF00237843

Bedell, H. E., Chung, S. T. L., & Patel, S. S. (2004). Attenuation of perceived motion smear during vergence and pursuit tracking. *Vision research*, *44*(9), 895–902. doi:10.1016/j.visres.2003.11.006

Bedell, H. E., Tong, J., & Aydin, M. (2010). The perception of motion smear during eye and head movements. *Vision research*, *50*(24), 2692–701. doi:10.1016/j.visres.2010.09.025

Bennett, M. R., & Hacker, P. M. S. (2001). Perception and memory in neuroscience: a conceptual analysis. *Progress in neurobiology*, *65*(6), 499–543. doi:10.1016/S0301-0082(01)00020-X

Bennett, S. J., & Barnes, G. R. (2004). Predictive smooth ocular pursuit during the transient disappearance of a visual target. *Journal of neurophysiology*, *92*(1), 578–90. doi:10.1152/jn.01188.2003

Berkels, B., Kondermann, C., Garbe, C., & Rumpf, M. (2009). Reconstructing Optical Flow Fields by Motion Inpainting. In D. Cremers (Ed.), *EMMCVPR '09 Proceedings of the 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition* (Vol. LNCS 5681, pp. 388–400). doi:10.1007/978-3-642-03641-5_29

Berntsen, D., & Bohn, A. (2010). Remembering and forecasting: The relation between autobiographical memory and episodic future thinking. *Memory & cognition*, *38*(3), 265–78. doi:10.3758/MC.38.3.265

Bilder, C. R. (2009). Human or Cylon? Group testing on Battlestar Galactica. *Chance (New York, N.Y.)*, *22*(3), 46–50. doi:10.1007/s00144-009-0030-1

Billock, V. A. (1997). Very short-term visual memory via reverberation: a role for the cortico-thalamic excitatory circuit in temporal filling-in during blinks and saccades? *Vision research*, *37*(7), 949–53. doi:10.1016/S0042-6989(96)00257-X

Binda, P., Cicchini, G. M., Burr, D. C., & Morrone, M. C. (2009). Spatiotemporal distortions of visual perception at the time of saccades. *Journal of Neuroscience*, *29*(42), 13147–57. doi:10.1523/JNEUROSCI.3723-09.2009

Blake, R., & Lee, S.-H. (2005). The role of temporal structure in human vision. *Behavioral and cognitive neuroscience reviews*, *4*(1), 21–42. doi:10.1177/1534582305276839

Blakeslee, B., & McCourt, M. E. (2008). Nearly instantaneous brightness induction. *Journal of vision*, *8*(2), 15.1–8. doi:10.1167/8.2.15

Blohm, G., & Crawford, J. D. (2009). Fields of gain in the brain. *Neuron*, *64*(5), 598–600. doi:10.1016/j.neuron.2009.11.022

Bloomfield, S. A., & Völgyi, B. (2009). The diverse functional roles and regulation of neuronal gap junctions in the retina. *Nature reviews. Neuroscience*, *10*(7), 495–506. doi:10.1038/nrn2636

Blue, R. C., & Szirko, M. (2005). A visual yet non-optical subjective intonation: una entonación subjetiva visual pero no óptica. *Electroneurobiología*, *13*(3), 299–300. Retrieved from http://electroneubio.secyt.gov.ar/non-optic_visual_intonation.htm

Booch, G. (2006). The Accidental Architecture. *IEEE Software*, *23*(3), 9–11. doi:10.1109/MS.2006.86

Boring, E. G. (1936). Psychophysiological systems and isomorphic relations. *Psychological Review*, *43*(6), 565–587. doi:10.1037/h0059300

Borra, E., Ichinohe, N., Sato, T., Tanifuji, M., & Rockland, K. S. (2010). Cortical connections to area TE in monkey: hybrid modular and distributed organization. *Cerebral Cortex*, *20*(2), 257–70. doi:10.1093/cercor/bhp096

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of vision*, *11*(5), 4.1–34. doi:10.1167/11.5.4

Brady, T. F., Konkle, T., Oliva, A., & Alvarez, G. A. (2009). Detecting changes in real-world objects: The relationship between visual long-term memory and change blindness. *Communicative & integrative biology*, *2*(1), 1–3. doi:10.4161/cib.2.1.7297

Brainard, D. H., & Maloney, L. T. (2011). Surface color perception and equivalent illumination models. *Journal of vision*, *11*(5), 1.1–18. doi:10.1167/11.5.1

Brainard, D. H., Williams, D. R., & Hofer, H. J. (2008). Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots. *Journal of vision*, *8*(5), 15.1–23. doi:10.1167/8.5.15

Breakspear, M., Jirsa, V. K., & Deco, G. (2010). Computational models of the brain: from structure to function. *NeuroImage*, *52*(3), 727–30. doi:10.1016/j.neuroimage.2010.05.061

Breitmeyer, B. G., & Jacob, J. (2012). Microgenesis of surface completion in visual objects: Evidence for filling-out. *Vision research*, *55*, 11–18. doi:10.1016/j.visres.2011.12.010

Breitmeyer, B. G., Kropfl, W., & Julesz, B. (1982). The existence and role of retinotopic and spatiotopic forms of visual persistence. *Acta psychologica*, *52*(3), 175–96. doi:10.1016/0001-6918(82)90007-5

Breitmeyer, B. G., Öğmen, H., & Chen, J. (2004). Unconscious priming by color and form: different processes and levels. *Consciousness and cognition*, *13*(1), 138–57. doi:10.1016/j.concog.2003.07.004

Bridgeman, B. (2007). Efference copy and its limitations. *Computers in biology and medicine*, *37*(7), 924–9. doi:10.1016/j.compbiomed.2006.07.001

Bridgeman, B., & Tseng, P. (2011). Embodied cognition and the perception-action link. *Physics of life reviews*, *8*(1), 73–85. doi:10.1016/j.plrev.2011.01.002

Briggs, F., & Usrey, W. M. (2011). Corticogeniculate feedback and visual processing in the primate. *Journal of physiology*, *589*(1), 33–40. doi:10.1113/jphysiol.2010.193599

Brill, M. H., & Robertson, A. R. (2007). Open Problems on the Validity of Grassmann's Laws. In J. Schanda (Ed.), *Colorimetry* (pp. 245–259). Hoboken, NJ, USA: John Wiley & Sons, Inc. doi:10.1002/9780470175637.ch10

Bristow, D., Frith, C., & Rees, G. (2005). Two distinct neural effects of blinking on human visual processing. *NeuroImage*, *27*(1), 136–45. doi:10.1016/j.neuroimage.2005.03.037

Bristow, D., Haynes, J.-D., Sylvester, R., Frith, C. D., & Rees, G. (2005). Blinking suppresses the neural response to unchanging retinal stimulation. *Current Biology*, *15*(14), 1296–300. doi:10.1016/j.cub.2005.06.025

Broackes, J. (1992). The Autonomy of Colour. In D. Charles & K. Lennon (Eds.), *Reduction, Explanation, and Realism* (pp. 421–465). Oxford: Clarendon Press. Retrieved from http://www.oup.com/us/catalog/general/subject/Philosophy/History/Modern/?view= usa&ci=9780198751311

Brockmeier, J. (2010). After the Archive: Remapping Memory. *Culture & Psychology*, *16*(1), 5–35. doi:10.1177/1354067X09353212

Brockmole, J. R., & Irwin, D. E. (2005). Eye movements and the integration of visual memory and visual perception. *Perception & psychophysics*, *67*(3), 495–512. doi:10.3758/BF03193327

Brooks, R. A. (1985). A Mobile Robot Project. Retrieved from http://hdl.handle.net/1721.1/41220

Brown, J. M. (2009). Visual streams and shifting attention. *Progress in Brain Research, Vol. 176: Attention* (Vol. 176, pp. 47–63). doi:10.1016/S0079-6123(09)17604-5

Brown, R. J., & Thurmond, J. B. (1993). Preattentive and cognitive effects on perceptual completion at the blind spot. *Perception & Psychophysics*, *53*(2), 200–209. doi:10.3758/BF03211730

Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, *53*(6), 453–462. doi:10.1016/j.jmp.2009.09.002

Browning, N. A., Grossberg, S., & Mingolla, E. (2009). A neural model of how the brain computes heading from optic flow in realistic scenes. *Cognitive psychology*, *59*(4), 320–56. doi:10.1016/j.cogpsych.2009.07.002

Bruce, C. J., & Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *Journal of neurophysiology*, *53*(3), 603–35. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3981231

Bubic, A., Von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in human neuroscience*, *4*(March), 25. doi:10.3389/fnhum.2010.00025

Bulakowski, P. F., Post, R. B., & Whitney, D. (2011). Reexamining the possible benefits of visual crowding: dissociating crowding from ensemble percepts. *Attention, perception & psychophysics*, *73*(4), 1003–9. doi:10.3758/s13414-010-0086-2

Burke, L. (1952). On the tunnel effect. *Quarterly Journal of Experimental Psychology*, *4*(3), 121–138. doi:10.1080/17470215208416611

Burr, D. C. (1980). Motion smear. *Nature*, *284*(5752), 164–5. doi:10.1038/284164a0

Burr, D. C. (2005). Vision: in the blink of an eye. *Current Biology*, *15*(14), R554–6. doi:10.1016/j.cub.2005.07.007

Burr, D. C., Cicchini, G. M., Binda, P., & Morrone, M. C. (2011). How transient "remapping" of neuronal receptive fields mediates perceptual stability. *Journal of Vision*, *11*(11), 537–537. doi:10.1167/11.11.537

Burr, D. C., & Morgan, M. J. (1997). Motion deblurring in human vision. *Proceedings. Biological sciences / The Royal Society*, *264*(1380), 431–6. doi:10.1098/rspb.1997.0061

Burr, D. C., & Morrone, M. C. (2005). Eye movements: building a stable world from glance to glance. *Current Biology*, *15*(20), R839–40. doi:10.1016/j.cub.2005.10.003

Burr, D. C., & Morrone, M. C. (2011). Spatiotopic coding and remapping in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1564), 504–15. doi:10.1098/rstb.2010.0244

Burr, D. C., & Thompson, P. (2011). Motion psychophysics: 1985-2010. *Vision research*, *51*(13), 1431–1456. doi:10.1016/j.visres.2011.02.008

Buschman, T. J., & Miller, E. K. (2010). Shifting the spotlight of attention: evidence for discrete computations in cognition. *Frontiers in human neuroscience*, *4*(November), 194. doi:10.3389/fnhum.2010.00194

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron,* *68*(3), 362–85. doi:10.1016/j.neuron.2010.09.023

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological review,* *114*(2), 340–75. doi:10.1037/0033-295X.114.2.340

Calkins, D. J. (2001). Seeing with S cones. *Progress in Retinal and Eye Research, 20*(3), 255–287. doi:10.1016/S1350-9462(00)00026-4

Capilla, P., Malo, J., Luque, M. J., & Artigas, J. M. (1998). Colour representation spaces at different physiological levels: a comparative analysis. *Journal of Optics*, *29*(5), 324–338. doi:10.1088/0150-536X/29/5/003

Carandini, M. (2012). From circuits to behavior: a bridge too far? *Nature neuroscience,* *15*(4), 507–9. doi:10.1038/nn.3043

Cardoso-Leite, P., & Gorea, A. (2010). On the perceptual/motor dissociation: a review of concepts, theory, experimental paradigms and data interpretations. *Seeing and perceiving*, *23*(2), 89–151. doi:10.1163/187847510X503588

Castet, E. (2010). Perception of Intra-saccadic Motion. In U. J. Ilg & G. S. Masson (Eds.), *Dynamics of Visual Motion Processing* (pp. 213–238). Boston, MA: Springer US. doi:10.1007/978-1-4419-0781-3_10

Chang, S. W. C., & Snyder, L. H. (2010). Idiosyncratic and systematic aspects of spatial representations in the macaque parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(17), 1–6. doi:10.1073/pnas.0913209107

Chessa, M., Sabatini, S. P., & Solari, F. (2009). A Fast Joint Bioinspired Algorithm for Optic Flow and Two-Dimensional Disparity Estimation. In M. Fritz, B. Schiele, & J. H. Piater (Eds.), *Computer Vision Systems* (Vol. 5815, pp. 184–193). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-642-04667-4_19

Chicharro, D., & Ledberg, A. (2012). When two become one: the limits of causality analysis of brain dynamics. (T. Wennekers, Ed.)*PloS one, 7*(3), e32466. doi:10.1371/journal.pone.0032466

Chua, F. K. (2011). A new object captures attention--but only when you know which objects are old. *Attention, perception & psychophysics*, *73*(3), 797–808. doi:10.3758/s13414-010-0066-6

Ciaramitaro, V. M., Mitchell, J. F., Stoner, G. R., Reynolds, J. H., & Boynton, G. M. (2011). Object-based attention to one of two superimposed surfaces alters responses

in human early visual cortex. *Journal of neurophysiology*, *105*(3), 1258–1265. doi:10.1152/jn.00680.2010

Cohen, M. A., Alvarez, G. A., & Nakayama, K. (2011). Natural-scene perception requires attention. *Psychological science*, *22*(9), 1165–72. doi:10.1177/0956797611419168

Cohen, M. A., & Grossberg, S. (1984). Neural dynamics of brightness perception: features, boundaries, diffusion, and resonance. *Perception & psychophysics*, *36*(5), 428–56. doi:10.1016/S0166-4115(08)61762-8

Colby, C. L., Berman, R. A., Heiser, L. M., & Saunders, R. C. (2005). Corollary discharge and spatial updating: when the brain is split, is space still unified? *Progress in Brain Research, Vol. 149: Cortical Function: a View from the Thalamus* (Vol. 149, pp. 187–205). doi:10.1016/S0079-6123(05)49014-7

Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1995). Oculocentric spatial representation in parietal cortex. *Cerebral Cortex*, *5*(5), 470–81. doi:10.1093/cercor/5.5.470

Colliaux, D., Molter, C., & Yamaguchi, Y. (2009). Working memory dynamics and spontaneous activity in a flip-flop oscillations network model with a Milnor attractor. *Cognitive neurodynamics*, *3*(2), 141–51. doi:10.1007/s11571-009-9078-0

Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & psychophysics*, *27*(3), 183–228. doi:10.3758/BF03204258

Compte, A., & Wang, X.-J. (2006). Tuning curve shift by attention modulation in cortical neurons: a computational study of its mechanisms. *Cerebral Cortex*, *16*(6), 761–78. doi:10.1093/cercor/bhj021

Connor, C. E. (2001). Shifting receptive fields. *Neuron*, *29*(3), 548–9. doi:10.1016/S0896-6273(01)00228-8

Coppola, D. M., & Purves, D. (1996). The extraordinarily rapid disappearance of entoptic images. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(15), 8001–4. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8755592

Cornelissen, F. W., Wade, A. R., Vladusich, T., Dougherty, R. F., & Wandell, B. a. (2006). No functional magnetic resonance imaging evidence for brightness and color filling-in in early human visual cortex. *Journal of Neuroscience*, *26*(14), 3634–41. doi:10.1523/JNEUROSCI.4382-05.2006

Cosmides, L., & Tooby, J. (1994). Beyond intuition and instinct blindness: toward an evolutionarily rigorous cognitive science. *Cognition*, *50*(1-3), 41–77. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8039372

Craddock, M., Martinovic, J., & Lawson, R. (2011). An advantage for active versus passive aperture-viewing in visual object recognition. *Perception*, *40*(10), 1154–63. doi:10.1068/p6974

Crapse, T. B., & Sommer, M. A. (2008). Corollary discharge circuits in the primate brain. *Current opinion in neurobiology*, *18*(6), 552–7. doi:10.1016/j.conb.2008.09.017

Crespi, S., Biagi, L., D'Avossa, G., Burr, D. C., Tosetti, M., & Morrone, M. C. (2011). Spatiotopic Coding of BOLD Signal in Human Visual Cortex Depends on Spatial Attention. *PloS one*, *6*(7), e21661. doi:10.1371/journal.pone.0021661

Crook, J. D., Manookin, M. B., Packer, O. S., & Dacey, D. M. (2011). Horizontal cell feedback without cone type-selective inhibition mediates "red-green" color opponency in midget ganglion cells of the primate retina. *Journal of neuroscience : the official journal of the Society for Neuroscience*, *31*(5), 1762–72. doi:10.1523/JNEUROSCI.4385-10.2011

Crossland, M. D., & Bex, P. J. (2009). Spatial alignment over retinal scotomas. *Investigative ophthalmology & visual science*, *50*(3), 1464–9. doi:10.1167/iovs.08-2690

Crouzet, S. M., Kirchner, H., & Thorpe, S. J. (2010). Fast saccades toward faces: face detection in just 100 ms. *Journal of vision*, *10*(4), 16.1–17. doi:10.1167/10.4.16

Cudeiro, J. (2008). [The visual world: perceiving the future]. *Archivos de la Sociedad Española de Oftalmología*, *83*(11), 635–6. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/19006014

Cullen, K. E. (2004). Sensory signals during active versus passive movement. *Current opinion in neurobiology*, *14*(6), 698–706. doi:10.1016/j.conb.2004.10.002

Cunningham, D. W. (2008). Visual prediction as indicated by perceptual adaptation to temporal delays and discrete stimulation. *Behavioral and Brain Sciences*, *31*(02). doi:10.1017/S0140525X08003865

Cuntz, H. (2012). The dendritic density field of a cortical pyramidal cell. *Frontiers in neuroanatomy*, *6*(February), 2. doi:10.3389/fnana.2012.00002

Curcio, C. A., Allen, K. A., Sloan, K. R., Lerea, C. L., Hurley, J. B., Klock, I. B., & Milam, A. H. (1991). Distribution and morphology of human cone photoreceptors stained with anti-blue opsin. *Journal of comparative neurology*, *312*(4), 610–24. doi:10.1002/cne.903120411

Curcio, C. A., Sloan, K. R., Kalina, R. E., & Hendrickson, A. E. (1990). Human photoreceptor topography. *Journal of comparative neurology*, *292*(4), 497–523. doi:10.1002/cne.902920402

Cutsuridis, V. (2010). Neural accumulator models of decision making in eye movements. In A. Hussain, I. Aleksander, L. S. Smith, A. K. Barros, R. Chrisley, & V. Cutsuridis (Eds.), *Advances in experimental medicine and biology* (Vol. 657, pp. 61–72). New York, NY: Springer. doi:10.1007/978-0-387-79100-5_4

Dacey, D. M. (1999). Primate retina: cell types, circuits and color opponency. *Progress in retinal and eye research*, *18*(6), 737–63. doi:10.1016/S1350-9462(98)00013-5

Daunizeau, J., David, O., & Stephan, K. E. (2009). Dynamic causal modelling: A critical review of the biophysical and statistical foundations. *NeuroImage*, *58*(2), 312–322. doi:10.1016/j.neuroimage.2009.11.062

Davidson, P. R., & Wolpert, D. M. (2005). Widespread access to predictive models in the motor system: a short review. *Journal of neural engineering*, *2*(3), S313–9. doi:10.1088/1741-2560/2/3/S11

Davies, P. (1995). Visual scotoma and visual afterimages: some evidence that the perceived visual afterimage may not be a purely retinal phenomenon: a single case study. *Perceptual and Motor Skills*, *81*(3), 849–850. doi:10.2466/pms.1995.81.3.849

De Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and cognition*, *18*(3), 569–77. doi:10.1016/j.concog.2009.03.002

De Graaf, T. A., Goebel, R., & Sack, A. T. (2011). Feedforward and quick recurrent processes in early visual cortex revealed by TMS? *NeuroImage*, *61*(3), 651–659. doi:10.1016/j.neuroimage.2011.10.020

De Hemptinne, C., Lefèvre, P., & Missal, M. (2008). Neuronal bases of directional expectation and anticipatory pursuit. *Journal of Neuroscience*, *28*(17), 4298–310. doi:10.1523/JNEUROSCI.5678-07.2008

De Kleer, J., & Brown, J. S. (1984). A qualitative physics based on confluences. *Artificial Intelligence*, *24*(1-3), 7–83. doi:10.1016/0004-3702(84)90037-7

De Lange, F. P., Jensen, O., & Dehaene, S. (2010). Accumulation of evidence during sequential decision making: the importance of top-down factors. *Journal of Neuroscience*, *30*(2), 731–8. doi:10.1523/JNEUROSCI.4080-09.2010

De Oliveira, R. F., Damisch, L., Hossner, E.-J., Oudejans, R. R. D., Raab, M., Volz, K. G., & Williams, M. A. (2009). The bidirectional links between decision making, perception, and action. *Progress in Brain Research, Vol. 174: Mind and Motion:*

The Bidirectional Link between Thought and Action (Vol. 174, pp. 85–93). doi:10.1016/S0079-6123(09)01308-9

De Weerd, P. (2006). Perceptual filling-in: More than the eye can see. *Progress in Brain Research, Vol. 154: Visual Perception, Part 1: Fundamentals of Vision: Low and Mid-Level Processes in Perception* (Vol. 154, pp. 227–245). doi:10.1016/S0079-6123(06)54012-9

De Weerd, P. (2011). Brightness Induction: A Study of Correlates of Surface Brightness in Early Visual Areas. *Chinese Journal of Psychology*, *53*(4), 429–435. Retrieved from http://www.airitilibrary.com/searchdetail.aspx?DocIDs=10139656-201112-201201100501-201201100501-429-435

Deering, M. F. (2005). A photon accurate model of the human eye. *ACM Transactions on Graphics*, *24*(3), 649. doi:10.1145/1073204.1073243

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS biology*, *5*(10), e260. doi:10.1371/journal.pbio.0050260

Delvenne, J.-F., & Holt, J. L. (2012). Splitting attention across the two visual fields in visual short-term memory. *Cognition*, *122*(2), 258–63. doi:10.1016/j.cognition.2011.10.015

Demb, J. B., & Brainard, D. H. (2010). Vision: Neurons show their true colours. *Nature*, *467*(7316), 670–1. doi:10.1038/467670b

Dennett, D. C. (1991). *Consciousness Explained* (p. 528). Retrieved from http://www.amazon.com/Consciousness-Explained-Daniel-C-Dennett/dp/0316180653

Dennett, D. C. (1998). No bridge over the stream of consciousness. *Behavioral and Brain Sciences*, *21*(6), 753–754. doi:10.1017/S0140525X98281751

Dennett, D. C. (2001). Surprise, surprise. *Behavioral and Brain Sciences*, *24*(5), 982–982. doi:10.1017/S0140525X01320113

Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences*, *15*(2), 183–201. doi:10.1017/S0140525X00068229

Destexhe, A. (2010). Dendrites do it in sequences. *Science*, *329*(5999), 1611–2. doi:10.1126/science.1196743

Deubel, H., Bridgeman, B., & Schneider, W. X. (2004). Different effects of eyelid blinks and target blanking on saccadic suppression of displacement. *Perception & psychophysics*, *66*(5), 772–8. doi:10.3758/BF03194971

Deubel, H., Schneider, W. X., & Bridgeman, B. (1996). Postsaccadic target blanking prevents saccadic suppression of image displacement. *Vision research*, *36*(7), 985–96. doi:10.1016/0042-6989(95)00203-0

Di Lollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology: General*, *109*(1), 75–97. doi:10.1037/0096-3445.109.1.75

Di Lollo, V. (1983). Icons no, iconic memory yes. *Behavioral and Brain Sciences*, *6*(01), 19. doi:10.1017/S0140525X00014370

Di Lollo, V., Clark, C. D., & Hogben, J. H. (1988). Separating visible persistence from retinal afterimages. *Perception & Psychophysics*, *44*(4), 363–368. doi:10.3758/BF03210418

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: the psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General*, *129*(4), 481–507. doi:10.1037/0096-3445.129.4.481

Di Lollo, V., Hogben, J. H., & Dixon, P. (1994). Temporal integration and segregation of brief visual stimuli: patterns of correlation in time. *Perception & psychophysics*, *55*(4), 373–86. doi:10.3758/BF03205295

Dickinson, C. A., & Intraub, H. (2008). Transsaccadic representation of layout: what is the time course of boundary extension? *Journal of Experimental Psychology: Human Perception and Performance*, *34*(3), 543–55. doi:10.1037/0096-1523.34.3.543

Ditchburn, R. W., & Ginsborg, B. L. (1952). Vision with a Stabilized Retinal Image. *Nature*, *170*(4314), 36–37. doi:10.1038/170036a0

Dixon, P., & Di Lollo, V. (1994). Beyond Visible Persistence: An Alternative Account of Temporal Integration and Segregation in Visual Processing. *Cognitive Psychology*, *26*(1), 33–63. doi:10.1006/cogp.1994.1002

Dodge, R. (1900). Visual perception during eye movement. *Psychological Review*, *7*(5), 454–465. doi:10.1037/h0067215

Dodgson, N. A. (2004). Variation and extrema of human interpupillary distance. *Proceedings of SPIE*, *5291*(January), 36–46. doi:10.1117/12.529999

Doi, E., Inui, T., Lee, T.-W., Wachtler, T., & Sejnowski, T. J. (2003). Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes. *Neural computation*, *15*(2), 397–417. doi:10.1162/089976603762552960

Dombeck, D. A., & Reiser, M. B. (2012). Real neuroscience in virtual worlds. *Current opinion in neurobiology*, *22*(1), 3–10. doi:10.1016/j.conb.2011.10.015

Domingos, P. (1999). The Role of Occam ' s Razor in Knowledge Discovery. *Data Mining and Knowledge Discovery*, *3*(4), 409–425. doi:10.1023/A:1009868929893

Donner, K., & Hemilä, S. (2007). Modelling the effect of microsaccades on retinal responses to stationary contrast patterns. *Vision research*, *47*(9), 1166–77. doi:10.1016/j.visres.2006.11.024

Dornhoefer, S. M., Unema, P. J. A., & Velichkovsky, B. M. (2002). Blinks, blanks and saccades: how blind we really are for relevant visual events. *Progress in Brain Research, Vol. 140: The Brain's eye: Neurobiological and clinical aspects of oculomotor research* (Vol. 140, pp. 119–131). doi:10.1016/S0079-6123(02)40046-5

Douglas, R. J., & Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology*, *17*(13), R496–500. doi:10.1016/j.cub.2007.04.024

Doya, K. (2008). Modulators of decision making. *Nature neuroscience*, *11*(4), 410–6. doi:10.1038/nn2077

Dresner, E. (2012). Turing, Matthews and Millikan: Effective Memory, Dispositionalism and Pushmepullyou Mental States. *International Journal of Philosophical Studies*, *20*(4), 461–472. doi:10.1080/09672559.2012.713375

Driver, J., & Baylis, G. C. (1989). Movement and visual attention: the spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(3), 448–56. doi:10.1037/0096-1523.15.3.448

Drummond, L., & Shomstein, S. (2010). Object-based attention: shifting or uncertainty? *Attention, perception & psychophysics*, *72*(7), 1743–55. doi:10.3758/APP.72.7.1743

Dudai, Y. (2009). Predicting not to predict too much: how the cellular machinery of memory anticipates the uncertain future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1255–62. doi:10.1098/rstb.2008.0320

Duhamel, J.-R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, *255*(5040), 90–2. doi:10.1126/science.1553535

D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, *3*(10), 1662–72. doi:10.1364/JOSAA.3.001662

Echterhoff, G. (2011). The pervasiveness of memory. *Culture & Psychology*, *17*(1), 3–9. doi:10.1177/1354067X10388856

Eckhorn, R. (1999). Neural mechanisms of scene segmentation: recordings from the visual cortex suggest basic circuits for linking field models. *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, *10*(3), 464–79. doi:10.1109/72.761705

Eizenman, M., Hallett, P. E., & Frecker, R. C. (1985). Power spectra for ocular drift and tremor. *Vision research*, *25*(11), 1635–40. doi:10.1016/0042-6989(85)90134-8

Eliasmith, C. (2007). How to build a brain: from function to implementation. *Synthese*, *159*(3), 373–388. doi:10.1007/s11229-007-9235-0

Enns, J. T., Brehaut, J. C., & Shore, D. I. (1999). The duration of a brief event in the mind's eye. *Journal of general psychology*, *126*(4), 355–72. doi:10.1080/00221309909595371

Ermentrout, G. B., & Kleinfeld, D. (2001). Traveling Electrical Waves in Cortex. *Neuron*, *29*(1), 33–44. doi:10.1016/S0896-6273(01)00178-7

Evinger, C., Manning, K. A., Pellegrini, J. J., Basso, M. A., Powers, A. S., & Sibony, P. A. (1994). Not looking while leaping: the linkage of blinking and saccadic gaze shifts. *Experimental Brain Research*, *100*(2), 337–344. doi:10.1007/BF00227203

Fair, D. A., Cohen, A. L., Power, J. D., Dosenbach, N. U. F., Church, J. A., Miezin, F. M., Schlaggar, B. L., et al. (2009). Functional brain networks develop from a "local to distributed" organization. *PLoS computational biology*, *5*(5), e1000381. doi:10.1371/journal.pcbi.1000381

Fang, F., Boyaci, H., & Kersten, D. (2009). Border ownership selectivity in human early visual cortex and its modulation by attention. *Journal of Neuroscience*, *29*(2), 460–5. doi:10.1523/JNEUROSCI.4628-08.2009

Farrajota, M., Rodrigues, J. M. F., & Du Buf, J. M. H. (2011). Optical flow by multi-scale annotated keypoints: A biological approach. *Proceedings of the Internation Conference on Bio-inspired Systems and Signal Processing (BIOSIGNALS 2011)*, 307–315. Retrieved from http://hdl.handle.net/10400.1/891

Fazl, A., Grossberg, S., & Mingolla, E. (2009). View-invariant object category learning, recognition, and search: how spatial and object attention are coordinated using surface-based attentional shrouds. *Cognitive psychology*, *58*(1), 1–48. doi:10.1016/j.cogpsych.2008.05.001

Fei-Fei, L., Iyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of vision*, *7*(1), 10.1–29. doi:10.1167/7.1.10

Feitosa-Santana, C., D'Antona, A. D., & Shevell, S. K. (2011). What kinds of contours bound the reach of filled-in color? *Journal of vision*, *11*(2), 2.1–11. doi:10.1167/11.2.2

Feldman, J. A. (1985). Four frames suffice: A provisional model of vision and space. *Behavioral and Brain Sciences*, *8*(02), 265. doi:10.1017/S0140525X00020707

Feldman, J. A., & Ballard, D. H. (1982). Connectionist Models and Their Properties. *Cognitive Science*, *6*(3), 205–254. doi:10.1207/s15516709cog0603_1

Fendrich, R., Rieger, J. W., & Heinze, H.-J. (2005). The effect of retinal stabilization on anorthoscopic percepts under free-viewing conditions. *Vision research*, *45*(5), 567–82. doi:10.1016/j.visres.2004.09.025

Ferree, C. E., & Rand, G. (1919). Chromatic thresholds of sensation from center to periphery of the retina and their bearing on color theory: Part I. *Psychological Review*, *26*(1), 16–41. doi:10.1037/h0071940

Ffytche, D. H., & Zeki, S. M. (2011). The primary visual cortex, and feedback to it, are not necessary for conscious vision. *Brain : a journal of neurology*, *134*(Pt 1), 247–57. doi:10.1093/brain/awq305

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field". *Vision Research*, *33*(2), 173–193. doi:10.1016/0042-6989(93)90156-Q

Field, G. D., Gauthier, J. L., Sher, A., Greschner, M., Machado, T. A., Jepson, L. H., Shlens, J., et al. (2010). Functional connectivity in the retina at the resolution of photoreceptors. *Nature*, *467*(7316), 673–7. doi:10.1038/nature09424

Fisher, T. L. (2006). A spotlight on the searchlight. *Journal of Neuroscience*, *26*(32), 8219–20. doi:10.1523/JNEUROSCI.2217-06.2006

Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Current Opinion in Neurobiology*, *10*(4), 438–443. doi:10.1016/S0959-4388(00)00113-6

Flombaum, J. I., Kundey, S. M., Santos, L. R., & Scholl, B. J. (2004). Dynamic object individuation in rhesus macaques: a study of the tunnel effect. *Psychological science*, *15*(12), 795–800. doi:10.1111/j.0956-7976.2004.00758.x

Flombaum, J. I., Scholl, B. J., & Pylyshyn, Z. W. (2008). Attentional resources in visual tracking through occlusion: the high-beams effect. *Cognition*, *107*(3), 904–31. doi:10.1016/j.cognition.2007.12.015

Flombaum, J. I., Scholl, B. J., & Santos, L. R. (2009). Spatiotemporal priority as a fundamental principle of object persistence. In B. Hood & L. Santos (Eds.), *Origins of Object Knowledge* (pp. 135–164). Oxford University Press. Retrieved from http://www.oup.com/us/catalog/general/subject/Psychology/Developmental/?view= usa&ci=9780199216895

Fogarty, C., & Stern, J. A. (1989). Eye movements and blinks: their relationship to higher cognitive processes. *International Journal of Psychophysiology*, *8*(1), 35–42. doi:10.1016/0167-8760(89)90017-2

Forani Jr., M., Oliveira, L. De, Volchan, E., Pessoa, L., Gattass, R., & Rocha-Miranda, C. E. (2003). Completion Through a Permanent Scotoma: Fast Interpolation Across the Blind Spot and the Processing of Occlusion. In L. Pessoa & P. De Weerd (Eds.), *Filling-In: from perceptual completion to cortical reorganization* (pp. 177–186). Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780195140132.003.0009

Forbus, K. D. (2011). Qualitative modeling. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(4), 374–391. doi:10.1002/wcs.115

Foster, D. H. (2003). Does colour constancy exist? *Trends in cognitive sciences*, *7*(10), 439–43. doi:10.1016/j.tics.2003.08.002

Foster, D. H. (2011). Color constancy. *Vision research*, *51*(7), 674–700. doi:10.1016/j.visres.2010.09.006

Fracasso, A., Caramazza, A., & Melcher, D. (2010). Continuous perception of motion and shape across saccadic eye movements. *Journal of vision*, *10*(13), 14.1–17. doi:10.1167/10.13.14

Francis, G., & Ericson, J. (2004). Using afterimages to test neural mechanisms for perceptual filling-in. *Neural Networks*, *17*(5-6), 737–52. doi:10.1016/j.neunet.2004.01.007

Francis, G., & Grossberg, S. (1996). Cortical dynamics of boundary segmentation and reset: persistence, afterimages, and residual traces. *Perception*, *25*(5), 543–67. doi:10.1068/p250543

Francis, G., Grossberg, S., & Mingolla, E. (1994). Cortical dynamics of feature binding and reset: control of visual persistence. *Vision research*, *34*(8), 1089–104. doi:10.1016/0042-6989(94)90012-4

Franconeri, S. L., Jonathan, S. V, & Scimeca, J. M. (2010). Tracking Multiple Objects Is Limited Only by Object Spacing, Not by Speed, Time, or Capacity. *Psychological Science*, *21*(7), 920–925. doi:10.1177/0956797610373935

Franconeri, S. L., Pylyshyn, Z. W., & Scholl, B. J. (2012). A simple proximity heuristic allows tracking of multiple objects through occlusion. *Attention, perception & psychophysics*, *74*(4), 691–702. doi:10.3758/s13414-011-0265-9

Frechette, E. S., Sher, A., Grivich, M. I., Petrusca, D., Litke, A. M., & Chichilnisky, E. J. (2005). Fidelity of the ensemble code for visual motion in primate retina. *Journal of neurophysiology*, *94*(1), 119–35. doi:10.1152/jn.01175.2004

Fred, A. L. N., & Jain, A. K. (2005). Combining multiple clusterings using evidence accumulation. *IEEE transactions on pattern analysis and machine intelligence*, *27*(6), 835–50. doi:10.1109/TPAMI.2005.113

Freeman, J., & Simoncelli, E. P. (2011). Metamers of the ventral stream. *Nature Neuroscience*, *14*(9), 1195–1201. doi:10.1038/nn.2889

Freeman, W. J. (2007). The place of "codes" in nonlinear neurodynamics. *Progress in Brain Research, Vol. 165: Computational Neuroscience: Theoretical Insights into Brain Function* (Vol. 165, pp. 447–462). doi:10.1016/S0079-6123(06)65028-0

Friedman, H. S., Zhou, H., & Von der Heydt, R. (2003). The coding of uniform colour figures in monkey visual cortex. *Journal of physiology*, *548*(Pt 2), 593–613. doi:10.1113/jphysiol.2002.033555

Friston, K. J. (2009). Modalities, modes, and models in functional neuroimaging. *Science*, *326*(5951), 399–403. doi:10.1126/science.1174521

Fu, P., Chang, R., O'Beirne, K., & King, W. M. (2008). Predictive disjunctive pursuit of virtual images perceived to move in depth. *Progress in Brain Research, Vol. 171: Using Eye Movements as an Experimental Probe of Brain function - A Symposium in Honor of Jean Büttner-Ennever*, *171*, 451–7. doi:10.1016/S0079-6123(08)00666-3

Fu, Y. X., Shen, Y., & Dan, Y. (2001). Motion-induced perceptual extrapolation of blurred visual targets. *Journal of Neuroscience*, *21*(20), RC172. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11588202

Fukuda, T. (2007). Structural organization of the gap junction network in the cerebral cortex. *Neuroscientist*, *13*(3), 199–207. doi:10.1177/1073858406296760

Ganmor, E., Segev, R., & Schneidman, E. (2011). The architecture of functional interaction networks in the retina. *Journal of Neuroscience*, *31*(8), 3044–54. doi:10.1523/JNEUROSCI.3682-10.2011

Gao, T., Gao, Z., Li, J., Sun, Z., & Shen, M. (2011). The perceptual root of object-based storage: an interactive model of perception and visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(6), 1803–23. doi:10.1037/a0025637

García-Pérez, M. A., & Peli, E. (2001). Intrasaccadic perception. *Journal of Neuroscience*, *21*(18), 7313–22. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11549741

Gardner, J. L., Merriam, E. P., Movshon, J. A., & Heeger, D. J. (2008). Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. *Journal of Neuroscience*, *28*(15), 3988–99. doi:10.1523/JNEUROSCI.5476-07.2008

Garrigan, P., Ratliff, C. P., Klein, J. M., Sterling, P., Brainard, D. H., & Balasubramanian, V. (2010). Design of a trichromatic cone array. *PLoS computational biology*, *6*(2), e1000677. doi:10.1371/journal.pcbi.1000677

Garzón, F. C. (2008). Towards a General Theory of Antirepresentationalism. *British Journal for the Philosophy of Science*, *59*(3), 259–292. doi:10.1093/bjps/axl007

Gauch, A., & Kerzel, D. (2009). Contributions of visible persistence and perceptual set to the flash-lag effect: focusing on flash onset abolishes the illusion. *Vision research*, *49*(24), 2983–91. doi:10.1016/j.visres.2009.09.018

Gawne, T. J., & Martin, J. M. (2000). Activity of primate V1 cortical neurons during blinks. *Journal of neurophysiology*, *84*(5), 2691–4. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11068010

Gawne, T. J., & Martin, J. M. (2002). Responses of primate visual cortical neurons to stimuli presented by flash, saccade, blink, and external darkening. *Journal of neurophysiology*, *88*(5), 2178–86. doi:10.1152/jn.00151.200

Georgeson, M. A., & Georgeson, J. M. (1985). On seeing temporal gaps between gratings: A criterion problem for measurement of visible persistence. *Vision Research*, *25*(11), 1729–1733. doi:10.1016/0042-6989(85)90145-2

Gibson, J. J. (1950). *Perception of the Visual World* (p. 235). Cambridge, MA: Riverside Press. Retrieved from http://archive.org/details/perceptionofvisu00jame

Gigerenzer, G. (1998). Surrogates for Theories. *Theory & Psychology*, *8*(2), 195–204. doi:10.1177/0959354398082006

Gigerenzer, G. (2010). Personal Reflections on Theory and Psychology. *Theory & Psychology*, *20*(6), 733–743. doi:10.1177/0959354310378184

Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, *1*(1), 107–143. doi:10.1111/j.1756-8765.2008.01006.x

Golomb, J. D., & Kanwisher, N. (2011). Higher Level Visual Cortex Represents Retinotopic, Not Spatiotopic, Object Location. *Cerebral Cortex*. doi:10.1093/cercor/bhr357

González, E. G., Weinstock, M., & Steinbach, M. J. (2007). Peripheral fading with monocular and binocular viewing. *Vision research*, *47*(1), 136–44. doi:10.1016/j.visres.2006.09.013

Goodale, M. A. (2008). Action without perception in human vision. *Cognitive neuropsychology*, *25*(7-8), 891–919. doi:10.1080/02643290801961984

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20–25. doi:10.1016/0166-2236(92)90344-8

Goossens, H. H. L. M., & Van Opstal, A. J. (2010). Differential effects of reflex blinks on saccade perturbations in humans. *Journal of neurophysiology*, *103*(3), 1685–95. doi:10.1152/jn.00788.2009

Graham, N. V. (2011). Beyond multiple pattern analyzers modeled as linear filters (as classical V1 simple cells): Useful additions of the last 25 years. *Vision research*, *51*(13), 1397–1430. doi:10.1016/j.visres.2011.02.007

Graziano, M. S. A., & Sigman, M. (2008). The dynamics of sensory buffers: geometric, spatial, and experience-dependent shaping of iconic memory. *Journal of vision*, *8*(5), 9.1–13. doi:10.1167/8.5.9

Green, N., & Heekeren, H. R. (2009). Perceptual decision making: a bidirectional link between mind and motion. *Progress in Brain Research, Vol. 174: Mind and Motion: The Bidirectional Link between Thought and Action* (Vol. 174, pp. 207–18). doi:10.1016/S0079-6123(09)01317-X

Grinvald, A., Arieli, A., Tsodyks, M., & Kenet, T. (2003). Neuronal assemblies: single cortical neurons are obedient members of a huge orchestra. *Biopolymers*, *68*(3), 422–36. doi:10.1002/bip.10273

Grossberg, S. (1983). The quantized geometry of visual space: The coherent computation of depth, form, and lightness. *Behavioral and Brain Sciences*, *6*(04), 625. doi:10.1017/S0140525X00017763

Grossberg, S. (1987). Cortical dynamics of three-dimensional form, color, and brightness perception: I. Monocular theory. *Perception & Psychophysics*, *41*(2), 87–116. doi:10.3758/BF03204874

Grossberg, S. (1991). Why do parallel cortical systems exist for the perception of static form and moving form? *Perception & psychophysics*, *49*(2), 117–41. doi:10.3758/BF03205033

Grossberg, S. (1994). 3-D vision and figure-ground separation by visual cortex. *Perception & Psychophysics*, *55*(1), 48–121. doi:10.3758/BF03206880

Grossberg, S. (2000). The complementary brain: unifying brain dynamics and modularity. *Trends in Cognitive Sciences*, *4*(6), 233–246. doi:10.1016/S1364-6613(00)01464-9

Grossberg, S. (2003). Filling-In the Forms: Surface and Boundary Interactions in Visual Cortex. In L. Pessoa & P. De Weerd (Eds.), *Filling-In: from perceptual completion to cortical reorganization* (pp. 13–37). Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780195140132.003.0002

Grossberg, S., & Grunewald, A. (2002). Temporal dynamics of binocular disparity processing with corticogeniculate interactions. *Neural Networks*, *15*(2), 181–200. doi:10.1016/S0893-6080(01)00149-6

Grossberg, S., Léveillé, J., & Versace, M. (2011). How do object reference frames and motion vector decomposition emerge in laminar cortical circuits? *Attention, perception & psychophysics*, *73*(4), 1147–70. doi:10.3758/s13414-011-0095-9

Grossberg, S., & Mingolla, E. (1985a). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics*, *38*(2), 141–171. doi:10.3758/BF03198851

Grossberg, S., & Mingolla, E. (1985b). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, *92*(2), 173–211. doi:10.1037/0033-295X.92.2.173

Grossberg, S., & Mingolla, E. (1993). Neural dynamics of motion perception: direction fields, apertures, and resonant grouping. *Perception & psychophysics*, *53*(3), 243–78. doi:10.3758/BF03205182

Grossberg, S., & Rudd, M. E. (1989). A neural architecture for visual motion perception: Group and element apparent motion. *Neural Networks*, *2*(6), 421–450. doi:10.1016/0893-6080(89)90042-7

Grossberg, S., & Rudd, M. E. (1992). Cortical dynamics of visual motion perception: Short-range and long-range apparent motion. *Psychological Review*, *99*(1), 78–121. doi:10.1037/0033-295X.99.1.78

Grossberg, S., & Todorović, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: a unified model of classical and recent phenomena. *Perception & psychophysics*, *43*(3), 241–77. doi:10.3758/BF03207869

Gulledge, A. T., Carnevale, N. T., & Stuart, G. J. (2012). Electrical Advantages of Dendritic Spines. (H. D. Mansvelder, Ed.)*PLoS ONE*, *7*(4), e36007. doi:10.1371/journal.pone.0036007

Haazebroek, P., Van Dantzig, S., & Hommel, B. (2011). A computational model of perception and action for cognitive robotics. *Cognitive processing*, *12*(4), 355–65. doi:10.1007/s10339-011-0408-x

Haber, R. N. (1983). The impending demise of the icon: A critique of the concept of iconic storage in visual information processing. *Behavioral and Brain Sciences*, *6*(01), 1. doi:10.1017/S0140525X0001428X

Haber, R. N. (1985). An icon can have no worth in the real world: Comments on Loftus, Johnson, and Shimamura's "How much is an icon worth?" *Journal of Experimental Psychology: Human Perception and Performance*, *11*(3), 374–378. doi:10.1037//0096-1523.11.3.374

Hafed, Z. M., & Krauzlis, R. J. (2010). Interactions Between Perception and Smooth Pursuit Eye Movements. In U. J. Ilg & G. S. Masson (Eds.), *Dynamics of Visual Motion Processing* (pp. 189–211). Boston, MA: Springer US. doi:10.1007/978-1-4419-0781-3_9

Hagan, M. a, Dean, H. L., & Pesaran, B. (2012). Spike-field activity in parietal area LIP during coordinated reach and saccade movements. *Journal of neurophysiology*, *107*(5), 1275–90. doi:10.1152/jn.00867.2011

Hall, N. J., & Colby, C. L. (2011). Remapping for visual stability. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1564), 528–39. doi:10.1098/rstb.2010.0248

Hallett, P. E., & Lightstone, A. D. (1976). Saccadic eye movements towards stimuli triggered by prior saccades. *Vision Research*, *16*(1), 99–106. doi:10.1016/0042-6989(76)90083-3

Hamburger, K., Prior, H., Sarris, V., & Spillmann, L. (2006). Filling-in with colour: different modes of surface completion. *Vision research*, *46*(6-7), 1129–38. doi:10.1016/j.visres.2005.08.013

Hamker, F. H., & Ziesche, A. (2011). Computational mechanisms of predictive remapping and visual stability. *Journal of Vision*, *11*(11), 523–523. doi:10.1167/11.11.523

Hamker, F. H., Zirnsak, M., Calow, D., & Lappe, M. (2008). The peri-saccadic perception of objects and space. *PLoS computational biology*, *4*(2), e31. doi:10.1371/journal.pcbi.0040031

Hamker, F. H., Zirnsak, M., Ziesche, A., & Lappe, M. (2011). Computational models of spatial updating in peri-saccadic perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1564), 554–571. doi:10.1098/rstb.2010.0229

Hammett, S. T., Georgeson, M. A., & Gorea, A. (1998). Motion blur and motion sharpening: temporal smear and local contrast non-linearity. *Vision research*, *38*(14), 2099–108. doi:10.1016/S0042-6989(97)00430-6

Han, Y., & Massey, S. C. (2005). Electrical synapses in retinal ON cone bipolar cells: subtype-specific expression of connexins. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(37), 13313–8. doi:10.1073/pnas.0505067102

Hansen, T., Pracejus, L., & Gegenfurtner, K. R. (2009). Color perception in the intermediate periphery of the visual field. *Journal of vision*, *9*(4), 26.1–12. doi:10.1167/9.4.26

Hartridge, H. (1922). Visual acuity and the resolving power of the eye. *Journal of physiology*, *57*(1-2), 52–67. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/16993599

Hartridge, H. (1944). Visibility of Blue and Yellow. *Nature*, *153*(3895), 775–776. doi:10.1038/153775a0

Hartridge, H. (1945). The Change from Trichromatic to Dichromatic Vision in the Human Retina. *Nature*, *155*(3944), 657–662. doi:10.1038/155657a0

Hartridge, H. (1947). The Visual Perception of Fine Detail. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *232*(592), 519–671. doi:10.1098/rstb.1947.0004

Hayhoe, M. M., & Williams, D. R. (1984). Disappearance of afterimages at "impossible" locations in space. *Perception*, *13*(4), 455–9. doi:10.1068/p130455

Hazelhoff, F., & Wiersma, H. (1924). Die Wahrnehmungszeit. *Zeitschrift für Psychologie / Journal of Psychology*, *96*, 171–188. Retrieved from http://books.google.ca/books?lr=&id=zos2AQAAIAAJ&dq=editions%3AUCAL_B650708&q=Die+Wahrnehmungszeit#search_anchor

Hecht, S., & Shlaer, S. (1936). INTERMITTENT STIMULATION BY LIGHT : V. THE RELATION BETWEEN INTENSITY AND CRITICAL FREQUENCY FOR

DIFFERENT PARTS OF THE SPECTRUM. *Journal of general physiology*, *19*(6), 965–77. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/19872976

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature reviews. Neuroscience*, *9*(6), 467–79. doi:10.1038/nrn2374

Hegdé, J. (2008). Time course of visual perception: coarse-to-fine processing and beyond. *Progress in neurobiology*, *84*(4), 405–39. doi:10.1016/j.pneurobio.2007.09.001

Helmholtz, H. von. (1925). *Treatise on Physiological Optics*. (J. P. C. Southall, Trans.) (3rd ed., Vol. 3). Optical Society of America. Retrieved from http://poseidon.sunyopt.edu/BackusLab/Helmholtz/

Herrington, T. M., & Assad, J. A. (2009). Neural activity in the middle temporal area and lateral intraparietal area during endogenously cued shifts of attention. *Journal of Neuroscience*, *29*(45), 14160–76. doi:10.1523/JNEUROSCI.1916-09.2009

Herrington, T. M., & Assad, J. A. (2010). Temporal sequence of attentional modulation in the lateral intraparietal area and middle temporal area during rapid covert shifts of attention. *Journal of Neuroscience*, *30*(9), 3287–96. doi:10.1523/JNEUROSCI.6025-09.2010

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., & Heeger, D. J. (2010). When size matters: attention affects performance by contrast or response gain. *Nature neuroscience*, *13*(12), 1554–9. doi:10.1038/nn.2669

Hesselmann, G., Sadaghiani, S., Friston, K. J., & Kleinschmidt, A. (2010). Predictive coding or evidence accumulation? False inference and neuronal fluctuations. *PloS one*, *5*(3), e9926. doi:10.1371/journal.pone.0009926

Heywood, S., & Churcher, J. (1971). Eye movements and the afterimage—I. Tracking the afterimage. *Vision Research*, *11*(10), 1163–1168. doi:10.1016/0042-6989(71)90120-9

Higgins, J. S., Irwin, D. E., Wang, R. F., & Thomas, L. E. (2009). Visual direction constancy across eyeblinks. *Attention, perception & psychophysics*, *71*(7), 1607–17. doi:10.3758/APP.71.7.1607

Hinton, G. E. (2007). To recognize shapes, first learn to generate images. *Progress in Brain Research, Vol. 165: Computational Neuroscience: Theoretical Insights into Brain Function* (Vol. 165, pp. 535–547). doi:10.1016/S0079-6123(06)65034-6

Hirsch, J., & Curcio, C. A. (1989). The spatial resolution capacity of human foveal retina. *Vision research*, *29*(9), 1095–101. doi:10.1016/0042-6989(89)90058-8

Ho, T. C., Brown, S. D., & Serences, J. T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *Journal of Neuroscience*, *29*(27), 8675–87. doi:10.1523/JNEUROSCI.5984-08.2009

Hochberg, J. (2007). In the Mind's Eye. In H. A. Peterson, Mary A; Gillam, Barbara; Sedgwick (Ed.), *In the mind's eye: Julian Hochberg on the Perception of Pictures, Films, and the World* (pp. 70–99). Oxford University Press. Retrieved from http://www.oupcanada.com/catalog/9780195176919.html

Hofer, H. J., Carroll, J., Neitz, J., Neitz, M., & Williams, D. R. (2005). Organization of the human trichromatic cone mosaic. *Journal of Neuroscience*, *25*(42), 9669–79. doi:10.1523/JNEUROSCI.2414-05.2005

Hohwy, J. (2007). Functional integration and the mind. *Synthese*, *159*(3), 315–328. doi:10.1007/s11229-007-9240-3

Holcombe, A. O. (2009). Seeing slow and seeing fast: two limits on perception. *Trends in cognitive sciences*, *13*(5), 216–21. doi:10.1016/j.tics.2009.02.005

Holland, M. K., & Tarlow, G. (1975). Blinking and thinking. *Perceptual and motor skills*, *41*(2), 503–6. doi:10.2466/pms.1975.41.2.403

Hollingworth, A. (2006). Visual memory for natural scenes: Evidence from change detection and visual search. *Visual Cognition*, *14*(4-8), 781–807. doi:10.1080/13506280500193818

Hollingworth, A., Simons, D. J., & Franconeri, S. L. (2010). New objects do not capture attention without a sensory transient. *Attention, perception & psychophysics*, *72*(5), 1298–310. doi:10.3758/APP.72.5.1298

Horowitz, T. S., & Cohen, M. A. (2010). Direction information in multiple object tracking is limited by a graded resource. *Attention, perception & psychophysics*, *72*(7), 1765–75. doi:10.3758/APP.72.7.1765

Horwitz, B., & Husain, F. T. (2007). Simulation Frameworks for Large-Scale Brain Systems. In V. K. Jirsa & A. McIntosh (Eds.), *Handbook of Brain Connectivity* (pp. 275–302). Berlin, Heidelberg: Springer. doi:10.1007/978-3-540-71512-2_9

Hsieh, P.-J., & Tse, P. U. (2010). BOLD signal in both ipsilateral and contralateral retinotopic cortex modulates with perceptual fading. *PloS one*, *5*(3), e9638. doi:10.1371/journal.pone.0009638

Hu, D., Zhou, Z., & Wang, Z. (2011). Processing real-world imagery with FACADE-based approaches. *Frontiers of Electrical and Electronic Engineering in China*, *6*(1), 120–136. doi:10.1007/s11460-011-0133-3

Huang, X., & Paradiso, M. A. (2008). V1 response timing and surface filling-in. *Journal of neurophysiology*, *100*(1), 539–47. doi:10.1152/jn.00997.2007

Hubbard, T. L., Hutchison, J. L., & Courtney, J. R. (2010). Boundary extension: findings and theories. *Quarterly journal of experimental psychology (2006)*, *63*(8), 1467–94. doi:10.1080/17470210903511236

Hunt, J. L. (2003). The Roget Illusion, the anorthoscope and the persistence of vision. *American Journal of Physics*, *71*(8), 774. doi:10.1119/1.1575766

Hupé, J.-M., Bordier, C., & Dojat, M. (2012). A BOLD signature of eyeblinks in the visual cortex. *NeuroImage*, *61*(1), 149–161. doi:10.1016/j.neuroimage.2012.03.001

Hérault, J., & Durette, B. (2007). Modeling Visual Perception for Image Processing. In F. Sandoval, A. Prieto, J. Cabestany, & M. Graña (Eds.), *Computational and Ambient Intelligence* (Vol. 4507, pp. 662–675). Berlin, Heidelberg: Springer. doi:10.1007/978-3-540-73007-1_80

Iijima, T., Witter, M. P., Ichikawa, M., Tominaga, T., Kajiwara, R., & Matsumoto, G. (1996). Entorhinal-Hippocampal Interactions Revealed by Real-Time Imaging. *Science*, *272*(5265), 1176–1179. doi:10.1126/science.272.5265.1176

Ikkai, A., & Curtis, C. E. (2008). Cortical activity time locked to the shift and maintenance of spatial attention. *Cerebral Cortex*, *18*(6), 1384–94. doi:10.1093/cercor/bhm171

Ilg, U. J., & Thier, P. (2003). Visual tracking neurons in primate area MST are activated by smooth-pursuit eye movements of an "imaginary" target. *Journal of neurophysiology*, *90*(3), 1489–502. doi:10.1152/jn.00272.2003

Imamizu, H. (2010). Prediction of sensorimotor feedback from the efference copy of motor commands: A review of behavioral and functional neuroimaging studies. *Japanese Psychological Research*, *52*(2), 107–120. doi:10.1111/j.1468-5884.2010.00428.x

Intraub, H. (2012). Rethinking visual scene perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(1), 117–127. doi:10.1002/wcs.149

Intraub, H., & Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 179–187. doi:10.1037/0278-7393.15.2.179

Irwin, D. E. (1991). Information integration across saccadic eye movements. *Cognitive psychology*, *23*(3), 420–56. doi:10.1016/0010-0285(91)90015-G

Isaacson, J. S., & Scanziani, M. (2011). How Inhibition Shapes Cortical Activity. *Neuron*, *72*(2), 231–243. doi:10.1016/j.neuron.2011.09.027

Jensen, M. S., Yao, R., Street, W. N., & Simons, D. J. (2011). Change blindness and inattentional blindness. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(5), 529–546. doi:10.1002/wcs.130

Johnson, M. H., & Munakata, Y. (2005). Processes of change in brain and cognitive development. *Trends in cognitive sciences*, *9*(3), 152–8. doi:10.1016/j.tics.2005.01.009

Jones, M., & Love, B. C. (2011). Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and brain sciences*, *34*(4), 169–188. doi:10.1017/S0140525X10003134

Kagan, I., Gur, M., & Snodderly, D. M. (2008). Saccades and drifts differentially modulate neuronal activity in V1: effects of retinal image motion, position, and extraretinal influences. *Journal of vision*, *8*(14), 19.1–25. doi:10.1167/8.14.19

Kalar, D. J., Garrigan, P., Wickens, T. D., Hilger, J. D., & Kellman, P. J. (2010). A unified model of illusory and occluded contour interpolation. *Vision research*, *50*(3), 284–99. doi:10.1016/j.visres.2009.10.011

Kaminer, J., Powers, A. S., Horn, K. G., Hui, C., & Evinger, C. (2011). Characterizing the spontaneous blink generator: an animal model. *Journal of Neuroscience*, *31*(31), 11256–67. doi:10.1523/JNEUROSCI.6218-10.2011

Kanai, R., Carlson, T. A., Verstraten, F. A. J., & Walsh, V. Z. (2009). Perceived timing of new objects and feature changes. *Journal of vision*, *9*(7), 5.1–13. doi:10.1167/9.7.5

Kanai, R., Wu, D.-A., Verstraten, F. A. J., & Shimojo, S. (2006). Discrete color filling beyond luminance gaps along perceptual surfaces. *Journal of vision*, *6*(12), 1380–95. doi:10.1167/6.12.4

Kandil, F. I., & Lappe, M. (2007). Spatio-temporal interpolation is accomplished by binocular form and motion mechanisms. *PloS one*, *2*(2), e264. doi:10.1371/journal.pone.0000264

Karson, C. N., Berman, K. F., Donnelly, E. F., Mendelson, W. B., Kleinman, J. E., & Wyatt, R. J. (1981). Speaking, thinking, and blinking. *Psychiatry Research*, *5*(3), 243–246. doi:10.1016/0165-1781(81)90070-6

Kastner, S., & McMains, S. A. (2007). Out of the spotlight: face to face with attention. *Nature neuroscience*, *10*(11), 1344–5. doi:10.1038/nn1107-1344

Katsuyama, N., Usui, N., Nose, I., & Taira, M. (2011). Perception of object motion in three-dimensional space induced by cast shadows. *NeuroImage*, *54*(1), 485–94. doi:10.1016/j.neuroimage.2010.07.075

Kaul, C., & White, A. L. (2010). When feature-based attention distorts neural representations. *Journal of Neuroscience*, *30*(31), 10261–2. doi:10.1523/JNEUROSCI.2777-10.2010

Kawabata, N. (1982). Visual information processing at the blind spot. *Perceptual and Motor Skills*, *55*(1), 95–104. doi:10.2466/pms.1982.55.1.95

Kawabata, N. (1983). Global interactions in perceptual completion at the blind spot. *Vision Research*, *23*(3), 275–279. doi:10.1016/0042-6989(83)90116-5

Kawabata, N. (1984). Perception at the blind spot and similarity grouping. *Perception & Psychophysics*, *36*(2), 151–158. doi:10.3758/BF03202675

Kayser, C., Montemurro, M. A., Logothetis, N. K., & Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron*, *61*(4), 597–608. doi:10.1016/j.neuron.2009.01.008

Keane, B. P., & Pylyshyn, Z. W. (2006). Is motion extrapolation employed in multiple object tracking? Tracking as a low-level, non-predictive function. *Cognitive psychology*, *52*(4), 346–68. doi:10.1016/j.cogpsych.2005.12.001

Keil, M. S., Cristóbal, G., Hansen, T., & Neumann, H. (2005). Recovering real-world images from single-scale boundaries with a novel filling-in architecture. *Neural Networks*, *18*(10), 1319–31. doi:10.1016/j.neunet.2005.04.003

Keitel, C., Andersen, S. K., Quigley, C., & Müller, M. M. (2012). Independent Effects of Attentional Gain Control and Competitive Interactions on Visual Stimulus Processing. *Cerebral Cortex*, *In Press*. doi:10.1093/cercor/bhs084

Keith, G. P., Blohm, G., & Crawford, J. D. (2010). Influence of saccade efference copy on the spatiotemporal properties of remapping: a neural network study. *Journal of neurophysiology*, *103*(1), 117–39. doi:10.1152/jn.91191.2008

Kelley, T. A., Serences, J. T., Giesbrecht, B., & Yantis, S. (2008). Cortical mechanisms for shifting and holding visuospatial attention. *Cerebral Cortex*, *18*(1), 114–25. doi:10.1093/cercor/bhm036

Kellman, P. J., Yin, C., & Shipley, T. F. (1998). A common mechanism for illusory and occluded object completion. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(3), 859–69. doi:10.1037/0096-1523.24.3.859

Kelly, D. H. (1990). Moving gratings and microsaccades. *Journal of the Optical Society of America A*, *7*(12), 2237. doi:10.1364/JOSAA.7.002237

Kevrekidis, I. G., & Samaey, G. (2009). Equation-free multiscale computation: algorithms and applications. *Annual review of physical chemistry*, *60*, 321–44. doi:10.1146/annurev.physchem.59.032607.093610

Keysers, C., Xiao, D.-K., Foldiak, P., & Perrett, D. I. (2005). Out of sight but not out of mind: the neurophysiology of iconic memory in the superior temporal sulcus. *Cognitive neuropsychology*, *22*(3), 316–32. doi:10.1080/02643290442000103

Keysers, C., Xiao, D.-K., Földiák, P., & Perrett, D. I. (2001). The Speed of Sight. *Journal of Cognitive Neuroscience*, *13*(1), 90–101. doi:10.1162/089892901564199

Khayat, P. S., Spekreijse, H., & Roelfsema, P. R. (2004). Visual information transfer across eye movements in the monkey. *Vision research*, *44*(25), 2901–17. doi:10.1016/j.visres.2004.06.018

Kingdom, F. A. A. (2011). Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision research*, *51*(7), 652–673. doi:10.1016/j.visres.2010.09.012

Kinser, J. M. (1999). Foveation by a pulse-coupled neural network. *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, *10*(3), 621–5. doi:10.1109/72.761721

Kirchner, H., & Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vision research*, *46*(11), 1762–76. doi:10.1016/j.visres.2005.10.002

Kirlik, A. (2007). Ecological Resources for Modeling Interactive Behavior and Embedded Cognition. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 194–210). New York, NY: Oxford University Press. Retrieved from http://www.oup.com/us/catalog/general/subject/Psychology/CognitivePsychology/?view=usa&ci=9780195189193

Kiverstein, J., & Clark, A. (2009). Introduction: Mind Embodied, Embedded, Enacted: One Church or Many? *Topoi*, *28*(1), 1–7. doi:10.1007/s11245-008-9041-4

Klatzky, R. L., Wu, B., & Stetten, G. (2010). The disembodied eye: consequences of displacing perception from action. *Vision research*, *50*(24), 2618–26. doi:10.1016/j.visres.2010.08.025

Klein, S. B., Robertson, T. E., & Delton, A. W. (2010). Facing the future: memory as an evolved system for planning future acts. *Memory & cognition*, *38*(1), 13–22. doi:10.3758/MC.38.1.13

Kleinfeld, D., Raccuia-Behling, F., & Chiel, H. J. (1990). Circuits constructed from identified Aplysia neurons exhibit multiple patterns of persistent activity. *Biophysical journal*, *57*(4), 697–715. doi:10.1016/S0006-3495(90)82591-1

Klier, E. M., & Angelaki, D. E. (2008). Spatial updating and the maintenance of visual constancy. *Neuroscience*, *156*(4), 801–18. doi:10.1016/j.neuroscience.2008.07.079

Knutsen, P. M., & Ahissar, E. (2009). Orthogonal coding of object location. *Trends in neurosciences*, *32*(2), 101–9. doi:10.1016/j.tins.2008.10.002

Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., & Mrsic-Flogel, T. D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature*, *473*, 87–91. doi:10.1038/nature09880

Koene, R. A., Tijms, B., Van Hees, P., Postma, F., De Ridder, A., Ramakers, G. J. A., Van Pelt, J., et al. (2009). NETMORPH: a framework for the stochastic generation of large scale neuronal networks with realistic neuron morphologies. *Neuroinformatics*, *7*(3), 195–210. doi:10.1007/s12021-009-9052-3

Koivisto, M. (2012). Is reentry critical for visual awareness of object presence? *Vision research*, *63*, 43–49. doi:10.1016/j.visres.2012.05.001

Kokaram, A. C. (2004). On Missing Data Treatment for Degraded Video and Film Archives: A Survey and a New Bayesian Approach. *IEEE Transactions on Image Processing*, *13*(3), 397–415. doi:10.1109/TIP.2004.823815

Kokkinos, I., Deriche, R., Maragos, P., & Faugeras, O. (2004). A Biologically Motivated and Computationally Tractable Model of Low and Mid-Level Vision Tasks. In T. Pajdla & J. Matas (Eds.), *Computer Vision – ECCV 2004* (Vol. 3022, pp. 506–517). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-540-24671-8_40

Komatsu, H. (2006). The neural mechanisms of perceptual filling-in. *Nature reviews. Neuroscience*, *7*(3), 220–31. doi:10.1038/nrn1869

Komatsu, H. (2011). Bridging Gaps at V1: Neural Responses for Filling-in and Completion at the Blind Spot. *Chinese Journal of Psychology*, *53*(4), 413–420. Retrieved from http://www.airitilibrary.com/searchdetail.aspx?DocIDs=10139656-201112-201201100501-201201100501-413-420

Komatsu, H., Kinoshita, M., & Murakami, I. (2000). Neural responses in the retinotopic representation of the blind spot in the macaque V1 to stimuli for perceptual filling-in. *Journal of Neuroscience*, *20*(24), 9310–9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11125010

Komatsu, H., Kinoshita, M., & Murakami, I. (2002). Neural responses in the primary visual cortex of the monkey during perceptual filling-in at the blind spot. *Neuroscience research*, *44*(3), 231–6. doi:10.1016/S0168-0102(02)00149-9

Komatsu, H., & Murakami, I. (1994). Huge receptive fields of the neurons in the retinotopic map of the macaque V1 representing the blind spot. *Neuroscience Research Supplements*, *19*, S199. doi:10.1016/0921-8696(94)92830-4

Konkle, T., & Oliva, A. (2011). Canonical visual size for real-world objects. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(1), 23–37. doi:10.1037/a0020413

Koriat, A., & Goldsmith, M. (1996). Memory metaphors and the real-life/laboratory controversy: Correspondence versus storehouse conceptions of memory. *Behavioral and Brain Sciences*, *19*(02), 167. doi:10.1017/S0140525X00042114

Kourtzi, Z., & Huberle, E. (2005). Spatiotemporal characteristics of form analysis in the human visual cortex revealed by rapid event-related fMRI adaptation. *NeuroImage*, *28*(2), 440–52. doi:10.1016/j.neuroimage.2005.06.017

Krekelberg, B. (2010). Saccadic suppression. *Current Biology*, *20*(5), R228–9. doi:10.1016/j.cub.2009.12.018

Kusunoki, M., & Goldberg, M. E. (2003). The time course of perisaccadic receptive field shifts in the lateral intraparietal area of the monkey. *Journal of neurophysiology*, *89*(3), 1519–27. doi:10.1152/jn.00519.2002

Kveraga, K., Ghuman, A. S., & Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain and cognition*, *65*(2), 145–68. doi:10.1016/j.bandc.2007.06.007

König, A. (1894). Über den menschlichen Sehpurpur und seine Bedeutung für das Sehen. *Sitzungsberichte der königlich preussischen Akademie der Wissenschaften zu Berlin* (Vol. 2, pp. 577–598). Retrieved from http://books.google.ca/books?id=LfInAQAAIAAJ&q=Über+den+menschlichen+Sehpurpur+und+seine+Bedeutung+für+das+Sehen#search_anchor

Land, E. H. (1986). Recent advances in retinex theory. *Vision Research*, *26*(1), 7–21. doi:10.1016/0042-6989(86)90067-2

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America (1917-1983)*, *61*(1), 1–11. doi:10.1364/JOSA.61.000001

Large, M.-E., Culham, J., Kuchinad, A., Aldcroft, A., & Vilis, T. (2008). fMRI reveals greater within- than between-hemifield integration in the human lateral occipital cortex. *European journal of neuroscience*, *27*(12), 3299–309. doi:10.1111/j.1460-9568.2008.06270.x

Lawlor, M., Holtmann-Rice, D., Huggins, P., Ben-Shahar, O., & Zucker, S. W. (2009). Boundaries, shading, and border ownership: A cusp at their interaction. *Journal of physiology, Paris*, *103*(1-2), 18–36. doi:10.1016/j.jphysparis.2009.05.005

Lee, B. B. (2011). Visual pathways and psychophysical channels in the primate. *Journal of physiology*, *589*(1), 41–47. doi:10.1113/jphysiol.2010.192658

Lee, B. B., Shapley, R. M., Hawken, M. J., & Sun, H. (2012). Spatial distributions of cone inputs to cells of the parvocellular pathway investigated with cone-isolating gratings. *Journal of the Optical Society of America A*, *29*(2), A223–32. doi:10.1364/JOSAA.29.00A223

Lee, W.-C. A., & Reid, R. C. (2011). Specificity and randomness: structure-function relationships in neural circuits. *Current opinion in neurobiology*, *21*(5), 801–807. doi:10.1016/j.conb.2011.07.004

Leibovitz, D. P. (2012a). A Unified Cognitive Model of Visual Filling-In Based on an Emergic Network Architecture - Animated Test Results. *Emergic Approach*. Retrieved from http://emergic.upwize.com/?page_id=26

Leibovitz, D. P. (2012b). Emergic Network. *Emergic Approach*. Retrieved from http://emergic.upwize.com/?page_id=6

Leibovitz, D. P., & West, R. L. (2012a). Cognitive Re-Use via Emergic Networks. In N. Rußwinkel, U. Drewitz, & H. van Rijn (Eds.), *11th International Conference on Cognitive Modeling (ICCM 2012)* (pp. 72–73). Berlin, Germany: Universitätsverlag der TU Berlin. Retrieved from http://www.iccm2012.com/proceedings/papers/0013/index.html

Leibovitz, D. P., & West, R. L. (2012b). Cognitive Re-Use via Emergic Networks (Poster). *11th International Conference on Cognitive Modeling (ICCM 2012)* (Vol. In Press). Berlin, Germany. Retrieved from http://www.iccm2012.com/proceedings/papers/0013/index.html

Lennie, P. (2003). Receptive fields. *Current Biology*, *13*(6), R216–R219. doi:10.1016/S0960-9822(03)00153-2

Lim, H., & Choe, Y. (2008). Extrapolative delay compensation through facilitating synapses and its relation to the flash-lag effect. *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, *19*(10), 1678–88. doi:10.1109/TNN.2008.2001002

Lin, J. K., Grier, D. G., & Cowan, J. D. (1997). Faithful Representation of Separable Distributions. *Neural Computation*, *9*(6), 1305–1320. doi:10.1162/neco.1997.9.6.1305

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision research*, *49*(10), 1194–204. doi:10.1016/j.visres.2008.05.025

Lisberger, S. G., & Westbrook, L. E. (1985). Properties of visual inputs that initiate horizontal smooth pursuit eye movements in monkeys. *Journal of Neuroscience*, *5*(6), 1662–73. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/4009252

Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron*, *55*(2), 313–23. doi:10.1016/j.neuron.2007.06.030

Liu, T., & Mance, I. (2011). Constant spread of feature-based attention across the visual field. *Vision research*, *51*(1), 26–33. doi:10.1016/j.visres.2010.09.023

Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and feature-based attention. *Vision research*, *47*(1), 108–13. doi:10.1016/j.visres.2006.09.017

Liverence, B. M., & Scholl, B. J. (2011). Selective Attention Warps Spatial Representation: Parallel but Opposing Effects on Attended Versus Inhibited Objects. *Psychological Science*, *22*(12), 1600–1608. doi:10.1177/0956797611422543

Livitz, G., Yazdanbakhsh, A., Eskew Jr., R. T., & Mingolla, E. (2010). Producing non-Hering Hue combinations using complementary chromatic induction. *Journal of Vision*, *9*(8), 363–363. doi:10.1167/9.8.363

Loftus, G. R., & Irwin, D. E. (1998). On the Relations among Different Measures of Visible and Informational Persistence. *Cognitive Psychology*, *35*(2), 135–199. doi:10.1006/cogp.1998.0678

Logvinenko, A. D. (2001). High-spatial-frequency tritanopia: S-filling-in or S-filtering-out? *Perception*, *30*(2), 223–232. doi:10.1068/p2846

Logvinenko, A. D., & Hutchinson, S. J. (2006). An effect of luminance contrast on high-spatial-frequency tritanopia. *Perception*, *35*(2), 171–184. doi:10.1068/p5051

London, M., & Häusser, M. (2005). Dendritic computation. *Annual review of neuroscience*, *28*, 503–32. doi:10.1146/annurev.neuro.28.061604.135703

Long, G. M. (1980). Iconic memory: a review and critique of the study of short-term visual storage. *Psychological bulletin*, *88*(3), 785–820. doi:10.1037/0033-2909.88.3.785

Long, G. M. (1982). Persisting problems in persistence: A response to Bowling and Lovegrove. *Perception & Psychophysics*, *32*(2), 192–194. doi:10.3758/BF03204279

Lorenceau, J. (2010). From Moving Contours to Object Motion: Functional Networks for Visual Form/Motion Processing. In U. J. Ilg & G. S. Masson (Eds.), *Dynamics of Visual Motion Processing* (pp. 3–36). Boston, MA: Springer US. doi:10.1007/978-1-4419-0781-3_1

Loschky, L. C., McConkie, G., Yang, J., & Miller, M. (2005). The limits of visual resolution in natural scene viewing. *Visual Cognition*, *12*(6), 1057–1092. doi:10.1080/13506280444000652

Lou, L., & Chen, J. (2003). Attention and Blind-Spot Phenomenology. *Psyche*, *9*, 1–14. Retrieved from http://www.theassc.org/journal_psyche/archive/vol_9_2003

Luo, D.-G., Kefalov, V. J., & Yau, K.-W. (2008). Phototransduction in Rods and Cones. In A. I. Basbaum, A. Kaneko, G. M. Shepherd, G. Westheimer, T. D. Albright, R. H. Masland, P. Dallos, et al. (Eds.), *The Senses: A Comprehensive Reference* (pp. 1.269–301). Elsevier. doi:10.1016/B978-012370880-9.00258-9

Mack, A., & Clarke, J. (2012). Gist perception requires attention. *Visual Cognition*, *20*(3), 300–327. doi:10.1080/13506285.2012.666578

MacKay, D. M. (1958). Perceptual stability of a stroboscopically lit visual field containing self-luminous objects. *Nature*, *181*(4607), 507–8. doi:10.1038/181507a0

Maertens, M., & Pollmann, S. (2007). Illusory contours do not pass through the "blind spot". *Journal of cognitive neuroscience*, *19*(1), 91–101. doi:10.1162/jocn.2007.19.1.91

Maex, R., & Steuber, V. (2009). The first second: models of short-term memory traces in the brain. *Neural Networks*, *22*(8), 1105–12. doi:10.1016/j.neunet.2009.07.022

Magee, J. C. (2000). Dendritic integration of excitatory synaptic input. *Nature reviews. Neuroscience*, *1*(3), 181–90. doi:10.1038/35044552

Magnussen, S., Spillmann, L., Stürzel, F., & Werner, J. S. (2001). Filling-in of the foveal blue scotoma. *Vision research*, *41*(23), 2961–7. doi:10.1016/S0042-6989(01)00178-X

Magnussen, S., Spillmann, L., Stürzel, F., & Werner, J. S. (2004). Unveiling the foveal blue scotoma through an afterimage. *Vision Research*, *44*(4), 377–383. doi:10.1016/j.visres.2003.09.023

Marcus, D. S., & Van Essen, D. C. (2002). Scene segmentation and attention in primate cortical areas V1 and V2. *Journal of neurophysiology*, *88*(5), 2648–58. doi:10.1152/jn.00916.2001

Marotta, J. J., McKeeff, T. J., & Behrmann, M. (2003). Hemispatial neglect: its effects on visual perception and visually guided grasping. *Neuropsychologia*, *41*(9), 1262–1271. doi:10.1016/S0028-3932(03)00038-1

Martinez-Conde, S., Krauzlis, R. J., Miller, J. M., Morrone, M. C., Williams, D. R., & Kowler, E. (2008). Eye movements and the perception of a clear and stable visual world. *Journal of Vision*, *8*(14), i–1–1. doi:10.1167/8.14.i

Martinez-Conde, S., Macknik, S. L., & Hubel, D. H. (2004). The role of fixational eye movements in visual perception. *Nature reviews. Neuroscience*, *5*(3), 229–40. doi:10.1038/nrn1348

Martinez-Conde, S., Macknik, S. L., Troncoso, X. G., & Hubel, D. H. (2009). Microsaccades: a neurophysiological analysis. *Trends in neurosciences*, *32*(9), 463–75. doi:10.1016/j.tins.2009.05.006

Masson, G. S., Montagnini, A., & Ilg, U. J. (2010). When the Brain Meets the Eye: Tracking Object Motion. In U. J. Ilg & G. S. Masson (Eds.), *Dynamics of Visual Motion Processing* (pp. 161–188). Boston, MA: Springer US. doi:10.1007/978-1-4419-0781-3_8

Masson, G. S., & Stone, L. S. (2002). From following edges to pursuing objects. *Journal of neurophysiology*, *88*(5), 2869–73. doi:10.1152/jn.00987.2001

Mateeff, S., Popov, D., & Hohnsbein, J. (1993). Multi-aperture viewing: perception of figures through very small apertures. *Vision Research*, *33*(17), 2563–2567. doi:10.1016/0042-6989(93)90135-J

Mather, G., & Smith, D. R. R. (2002). Blur discrimination and its relation to blur-mediated depth perception. *Perception*, *31*(10), 1211–9. doi:10.1068/p3254

Mathôt, S., & Theeuwes, J. (2010). Evidence for the predictive remapping of visual attention. *Experimental Brain Research*, *200*(1), 117–22. doi:10.1007/s00221-009-2055-3

Matsumoto, M., & Komatsu, H. (2005). Neural responses in the macaque v1 to bar stimuli with various lengths presented on the blind spot. *Journal of neurophysiology*, *93*(5), 2374–87. doi:10.1152/jn.00811.2004

Maus, G. W., & Nijhawan, R. (2008). Motion extrapolation into the blind spot. *Psychological Science*, *19*(11), 1087–91. doi:10.1111/j.1467-9280.2008.02205.x

Mays, L. E., & Sparks, D. L. (1980). Saccades are spatially, not retinocentrically, coded. *Science*, *208*(4448), 1163–5. doi:10.1126/science.6769161

McAlpine, D., & Palmer, A. R. (2006). How General Are Neural Codes in Sensory Systems? In J. L. van Hemmen & T. J. Sejnowski (Eds.), *23 Problems in Systems Neuroscience* (pp. 283–303). Oxford University Press. doi:10.1093/acprof:oso/9780195148220.003.0014

McCamy, C. S. (2003). Colors of some small figures on colored grounds. *Color Research & Application*, *28*(4), 242–250. doi:10.1002/col.10159

McCloskey, M., & Watkins, M. J. (1978). The seeing-more-than-is-there phenomenon: implications for the locus of iconic storage. *Journal of Experimental Psychology: Human Perception and Performance*, *4*(4), 553–64. doi:10.1037/0096-1523.4.4.553

McCree, K. J. (1960). Small-field Tritanopia and the Effects of Voluntary Fixation. *Optica Acta: International Journal of Optics*, *7*(4), 317–323. doi:10.1080/713826349

McDermott, J. H. (2009). The cocktail party problem. *Current Biology*, *19*(22), R1024–7. doi:10.1016/j.cub.2009.09.005

McMains, S. A., & Kastner, S. (2011). Interactions of Top-Down and Bottom-Up Mechanisms in Human Visual Cortex. *Journal of Neuroscience*, *31*(2), 587–597. doi:10.1523/JNEUROSCI.3766-10.2011

McManus, J. N. J., Ullman, S., & Gilbert, C. D. (2008). A computational model of perceptual fill-in following retinal degeneration. *Journal of neurophysiology*, *99*(5), 2086–100. doi:10.1152/jn.00871.2007

Medendorp, W. P. (2011). Spatial constancy mechanisms in motor control. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1564), 476–91. doi:10.1098/rstb.2010.0089

Meeter, M., Van Der Stigchel, S., & Theeuwes, J. (2010). A competitive integration model of exogenous and endogenous eye movements. *Biological cybernetics*, *102*(4), 271–91. doi:10.1007/s00422-010-0365-y

Meijering, E. (2010). Neuron tracing in perspective. *Cytometry. Part A : the journal of the International Society for Analytical Cytology*, *77*(7), 693–704. doi:10.1002/cyto.a.20895

Melcher, D., & Colby, C. L. (2008). Trans-saccadic perception. *Trends in cognitive sciences*, *12*(12), 466–73. doi:10.1016/j.tics.2008.09.003

Memmesheimer, R.-M., & Timme, M. (2012). Non-additive coupling enables propagation of synchronous spiking activity in purely random networks. (L. J. Graham, Ed.)*PLoS computational biology*, *8*(4), e1002384. doi:10.1371/journal.pcbi.1002384

Merriam, E. P., Genovese, C. R., & Colby, C. L. (2007). Remapping in human visual cortex. *Journal of neurophysiology*, *97*(2), 1738–55. doi:10.1152/jn.00189.2006

Mitrani, L., & Dimitrov, G. (1978). Pursuit eye movements of a disappearing moving target. *Vision research*, *18*(5), 537–9. doi:10.1016/0042-6989(78)90199-2

Mitroff, S. R., & Scholl, B. J. (2004). Seeing the disappearance of unseen objects. *Perception*, *33*(10), 1267–1273. doi:10.1068/p5341no

Momiji, H., Bharath, A. A., Hankins, M. W., & Kennard, C. (2006). Numerical study of short-term afterimages and associate properties in foveal vision. *Vision research*, *46*(3), 365–81. doi:10.1016/j.visres.2005.08.031

Momiji, H., Hankins, M. W., Bharath, A. A., & Kennard, C. (2007). A numerical study of red-green colour opponent properties in the primate retina. *European journal of neuroscience*, *25*(4), 1155–65. doi:10.1111/j.1460-9568.2007.05362.x

Moore, C. M., & Enns, J. T. (2004). Object updating and the flash-lag effect. *Psychological Science*, *15*(12), 866–71. doi:10.1111/j.0956-7976.2004.00768.x

Moore, C. M., Mordkoff, J. T., & Enns, J. T. (2007). The path of least persistence: object status mediates visual updating. *Vision research*, *47*(12), 1624–30. doi:10.1016/j.visres.2007.01.030

Morgan, M. J., Findlay, J. M., & Watt, R. J. (1982). Aperture viewing: a review and a synthesis. *Quarterly journal of experimental psychology. A, Human experimental psychology*, *34*(Pt 2), 211–33. doi:10.1080/14640748208400837

Morgan, M. J., & Turnbull, D. F. (1978). Smooth eye tracking and the perception of motion in the absence of real movement. *Vision research*, *18*(8), 1053–9. doi:10.1016/0042-6989(78)90034-2

Morita, K., Okada, M., & Aihara, K. (2007). Selectivity and stability via dendritic nonlinearity. *Neural computation*, *19*(7), 1798–853. doi:10.1162/neco.2007.19.7.1798

Morris, A. P., Kubischik, M., Hoffmann, K.-P., Krekelberg, B., & Bremmer, F. (2012). Dynamics of eye-position signals in the dorsal visual system. *Current Biology*, *22*(3), 173–9. doi:10.1016/j.cub.2011.12.032

Muckli, L., Vetter, P., & Smith, F. (2011). Predictive coding - contextual processing in primary visual cortex V1. *Journal of Vision*, *11*(11), 25–25. doi:10.1167/11.11.25

Murakami, I. (1995). Motion aftereffect after monocular adaptation to filled-in motion at the blind spot. *Vision research*, *35*(8), 1041–5. doi:10.1016/0042-6989(94)00201-V

Murakami, I. (2006). Fixational eye movements and motion perception. *Progress in Brain Research, Vol. 154: Visual Perception, Part 1: Fundamentals of Vision: Low and Mid-Level Processes in Perception*, *514*, 193–209. doi:10.1016/S0079-6123(06)54010-5

Murray, M. M., Foxe, D. M., Javitt, D. C., & Foxe, J. J. (2004). Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. *Journal of Neuroscience*, *24*(31), 6898–903. doi:10.1523/JNEUROSCI.1996-04.2004

Nakayama, K. (1994). James J. Gibson: An appreciation. *Psychological Review*, *101*(2), 329–335. doi:10.1037/0033-295X.101.2.329

Nakayama, K., He, Z. J., & Shimojo, S. (1995). Visual surface representation: a critical link between lower-level and higher level vision. In S. M. Kosslyn & D. N. Osherson (Eds.), *In Invitation to Cognitive Science: Visual Cognition* (2nd ed., Vol. 2, pp. 1–70). MIT Press. Retrieved from http://visionlab.harvard.edu/Members/Ken/nakayamapub.htm

Nakayama, K., Shimojo, S., Anderson, B. L., Kramer, P., Bressan, P., & Spillmann, L. (2009). Nakayama, Shimojo, and Ramachandran's 1990 paper. *Perception*, *38*, 859–877. doi:10.1068/ldmk-nak

Namba, J., & Baldo, V. C. (2004). The modulation of the flash-lag effect by voluntary attention. *Perception*, *33*(5), 621–31. doi:10.1068/p5212

Needham, J. (1962). *Science and Civilisation in China: Vol. 4, Physics and Physical Technology; Part 1, Physics* (Vol. 4). Retrieved from http://www.cambridge.org/gb/knowledge/isbn/item1118936/?site_locale=en_GB

Neisser, U. (1967). *Cognitive Psychology* (p. 351). New York: Appleton-Century-Crofts. Retrieved from http://books.google.ca/books/about/Cognitive_psychology.html?id=gIJ9AAAAMAAJ

Neisser, U. (1976). *Cognition and Reality: Principles and Implications of Cognitive Psychology* (p. 230). W H Freeman. Retrieved from http://books.google.ca/books?id=28u7QgAACAAJ

Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., & Williams, D. R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, *35*(4), 783–92. doi:10.1016/S0896-6273(02)00818-8

Nemes, V. A., Whitaker, D., Heron, J., & McKeefry, D. J. (2011). Multiple spatial frequency channels in human visual perceptual memory. *Vision research*, *51*(23-24), 2331–2339. doi:10.1016/j.visres.2011.09.003

Newell, B. R., & Bröder, A. (2008). Cognitive processes, models and metaphors in decision research. *Judgment and Decision Making*, *3*(3), 195–204. Retrieved from http://journal.sjdm.org/vol3.3.htm

Niemeier, M., Crawford, J. D., & Tweed, D. B. (2003). Optimal transsaccadic integration explains distorted spatial perception. *Nature*, *422*(6927), 76–80. doi:10.1038/nature01439

Nijhawan, R. (1994). Motion extrapolation in catching. *Nature*, *370*(6487), 256–7. doi:10.1038/370256b0

Nijhawan, R. (2002). Neural delays, visual motion and the flash-lag effect. *Trends in cognitive sciences*, *6*(9), 387. doi:10.1016/S1364-6613(02)01963-0

Nijhawan, R. (2008). Visual prediction: psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, *31*(2), 179–98; discussion 198–239. doi:10.1017/S0140525X08003804

Nijhawan, R., & Wu, S. (2009). Compensating time delays with neural predictions: are predictions sensory or motor? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *367*(1891), 1063–78. doi:10.1098/rsta.2008.0270

Nishida, S. (2004). Motion-based analysis of spatial patterns by the human visual system. *Current Biology*, *14*(10), 830–9. doi:10.1016/j.cub.2004.04.044

Nishida, S. (2011). Advancement of motion psychophysics: Review 2001-2010. *Journal of vision*, *11*(5), 11.1–53. doi:10.1167/11.5.11

Norman, J. (2001). Perceptual Constancy: Direct versus Constructivist Theories. *International Encyclopedia of the Social & Behavioral Sciences* (pp. 11221–11224). Elsevier. doi:10.1016/B0-08-043076-7/01452-2

Oliva, A., & Schyns, P. G. (2000). Diagnostic colors mediate scene recognition. *Cognitive psychology*, *41*(2), 176–210. doi:10.1006/cogp.1999.0728

Oliva, A., & Torralba, A. (2006). Building the gist of a scene: the role of global image features in recognition. *Progress in brain research*, *155*, 23–36. doi:10.1016/S0079-6123(06)55002-2

Olshausen, B. A., & Anderson, C. H. (2010). Does the brain de-jitter retinal images? *Proceedings of the National Academy of Sciences of the United States of America*, *107*(46), 19607–8. doi:10.1073/pnas.1015709107

Olshausen, B. A., Anderson, C. H., & Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of

information. *Journal of Neuroscience, 13*(11), 4700–19. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8229193

Orabona, F., Metta, G., & Sandini, G. (2007). A Proto-object Based Visual Attention Model. *Attention in Cognitive Systems. Theories and Systems from an Interdisciplinary Viewpoint* (Vol. 4840, pp. 198–215). Springer. doi:10.1007/978-3-540-77343-6_13

Orban de Xivry, J.-J., Bennett, S. J., Lefèvre, P., & Barnes, G. R. (2006). Evidence for synergy between saccades and smooth pursuit during transient target disappearance. *Journal of neurophysiology*, *95*(1), 418–27. doi:10.1152/jn.00596.2005

Orban de Xivry, J.-J., Missal, M., & Lefèvre, P. (2008). A dynamic representation of target motion drives predictive smooth pursuit during target blanking. *Journal of vision*, *8*(15), 6.1–13. doi:10.1167/8.15.6

Orban de Xivry, J.-J., Missal, M., & Lefèvre, P. (2009). Smooth pursuit performance during target blanking does not influence the triggering of predictive saccades. *Journal of vision*, *9*(11), 7.1–16. doi:10.1167/9.11.7

Osborne, L. C., & Lisberger, S. G. (2009). Spatial and temporal integration of visual motion signals for smooth pursuit eye movements in monkeys. *Journal of neurophysiology*, *102*(4), 2013–25. doi:10.1152/jn.00611.2009

Otero-Millan, J., Troncoso, X. G., Macknik, S. L., Serrano-Pedraza, I., & Martinez-Conde, S. (2008). Saccades and microsaccades during visual fixation, exploration, and search: foundations for a common saccadic generator. *Journal of vision*, *8*(14), 21.1–18. doi:10.1167/8.14.21

O'Herron, P., & Von der Heydt, R. (2011). Representation of object continuity in the visual cortex. *Journal of vision*, *11*(2), 12.1–9. doi:10.1167/11.2.12

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, *24*(5), 939–73; discussion 973–1031. doi:10.1017/S0140525X01000115

Packer, O. S., & Williams, D. R. (1992). Blurring by fixational eye movements. *Vision research*, *32*(10), 1931–9. doi:10.1016/0042-6989(92)90052-K

Paffen, C. L. E., & Van der Stigchel, S. (2010). Shifting spatial attention makes you flip: Exogenous visual attention triggers perceptual alternations during binocular rivalry. *Attention, perception & psychophysics*, *72*(5), 1237–43. doi:10.3758/APP.72.5.1237

Page, M. (2000). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, *23*(4), 443–467. doi:10.1017/S0140525X00003356

Page, W. K., & Duffy, C. J. (2008). Cortical neuronal responses to optic flow are shaped by visual strategies for steering. *Cerebral Cortex*, *18*(4), 727–39. doi:10.1093/cercor/bhm109

Palmer, D. C., & Donahoe, J. W. (1992). Essentialism and selectionism in cognitive science and behavior analysis. *American psychologist*, *47*(11), 1344–58. doi:10.1037/0003-066X.47.11.1344

Palmer, E. M., Kellman, P. J., & Shipley, T. F. (2006). A theory of dynamic occluded and illusory object perception. *Journal of Experimental Psychology: General*, *135*(4), 513–41. doi:10.1037/0096-3445.135.4.513

Palmer, L., Murayama, M., & Larkum, M. (2012). Inhibitory Regulation of Dendritic Activity in vivo. *Frontiers in neural circuits*, *6*(May), 26. doi:10.3389/fncir.2012.00026

Panzeri, S., Brunel, N., Logothetis, N. K., & Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends in neurosciences*, *33*(3), 111–20. doi:10.1016/j.tins.2009.12.001

Paradiso, M. A., & Nakayama, K. (1991). Brightness perception and filling-in. *Vision research*, *31*(7-8), 1221–36. doi:10.1016/0042-6989(91)90047-9

Park, S., Intraub, H., Yi, D.-J., Widders, D., & Chun, M. M. (2007). Beyond the edges of a view: boundary extension in human scene-selective visual cortex. *Neuron*, *54*(2), 335–42. doi:10.1016/j.neuron.2007.04.006

Parker, D. (2010). Neuronal network analyses: premises, promises and uncertainties. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *365*(1551), 2315–28. doi:10.1098/rstb.2010.0043

Parks, N. A., & Corballis, P. M. (2010). Human transsaccadic visual processing: presaccadic remapping and postsaccadic updating. *Neuropsychologia*, *48*(12), 3451–8. doi:10.1016/j.neuropsychologia.2010.07.028

Parks, T. E. (1965). Post-Retinal Visual Storage. *American Journal of Psychology*, *78*(1), 145. doi:10.2307/1421101

Pastukhov, A., & Braun, J. (2010). Rare but precious: microsaccades are highly informative about attentional allocation. *Vision research*, *50*(12), 1173–84. doi:10.1016/j.visres.2010.04.007

Pelz, J. B., & Hayhoe, M. M. (1995). The role of exocentric reference frames in the perception of visual direction. *Vision research*, *35*(16), 2267–75. doi:10.1016/0042-6989(94)00314-C

Perlovsky, L. I. (2007). Neural Modeling Fields, Evolution of Consciousness and Cultures. *2007 International Joint Conference on Neural Networks*, 460–465. doi:10.1109/IJCNN.2007.4371000

Perrett, D. I., Xiao, D., Barraclough, N. E., Keysers, C., & Oram, M. W. (2009). Seeing the future: Natural image sequences produce "anticipatory" neuronal activity and bias perceptual report. *Quarterly Journal of Experimental Psychology*, *62*(11), 2081–104. doi:10.1080/17470210902959279

Perrone, J. A., & Krauzlis, R. J. (2008). Vector subtraction using visual and extraretinal motion signals: a new look at efference copy and corollary discharge theories. *Journal of vision*, *8*(14), 24.1–14. doi:10.1167/8.14.24

Pessoa, L., Thompson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, *21*(6), 723–48; discussion 748–802. doi:10.1017/S0140525X98001757

Petrov, Y., Popple, A. V, & McKee, S. P. (2007). Crowding and surround suppression: not to be confused. *Journal of vision*, *7*(2), 12.1–9. doi:10.1167/7.2.12

Picard, R. W. (1995). Light-years from Lena: video and image libraries of the future. *Proceedings., International Conference on Image Processing* (Vol. 1, pp. 310–313). IEEE Comput. Soc. Press. doi:10.1109/ICIP.1995.529708

Pikler, J. (1917). *Sinnesphysiologische Untersuchungen* (p. 532). Leipzig: Barth. Retrieved from http://books.google.ca/books?id=P9eQpe0d0d0C

Ploran, E. J., Nelson, S. M., Velanova, K., Donaldson, D. I., Petersen, S. E., & Wheeler, M. E. (2007). Evidence accumulation and the moment of recognition: dissociating perceptual recognition processes using fMRI. *Journal of Neuroscience*, *27*(44), 11912–24. doi:10.1523/JNEUROSCI.3522-07.2007

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, *317*(6035), 314–319. doi:10.1038/317314a0

Pola, J., & Wyatt, H. J. (1997). Offset dynamics of human smooth pursuit eye movements: effects of target presence and subject attention. *Vision research*, *37*(18), 2579–95. doi:10.1016/S0042-6989(97)00058-8

Poletti, M., Listorti, C., & Rucci, M. (2010). Stability of the visual world during eye drift. *Journal of Neuroscience*, *30*(33), 11143–50. doi:10.1523/JNEUROSCI.1925-10.2010

Pollack, J. B. (1993). On wings of knowledge: a review of Allen Newell's Unified Theories of Cognition. *Artificial Intelligence*, *59*(1-2), 355–369. doi:10.1016/0004-3702(93)90200-U

Polsky, A., Mel, B. W., & Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nature neuroscience*, *7*(6), 621–7. doi:10.1038/nn1253

Ponder, E., & Kennedy, W. P. (1927). On the act of blinking. *Quarterly Journal of Experimental Physiology*, *18*, 89–110. Retrieved from http://ep.physoc.org/content/18/2/89

Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning & Memory*, *2*(5), 509–522. doi:10.1037/0278-7393.2.5.509

Potter, M. C. (2012). Recognition and memory for briefly presented scenes. *Frontiers in psychology*, *3*(February), 32. doi:10.3389/fpsyg.2012.00032

Pouget, A., Ringach, D. L., & Landy, M. (2010). Data rules; but theory understands: An introduction to a Special Issue on "Mathematical Models of Visual Coding". *Vision Research*, *50*(22), 2189–2189. doi:10.1016/j.visres.2010.10.001

Poulet, J. F. A., & Hedwig, B. (2007). New insights into corollary discharges mediated by identified neural pathways. *Trends in neurosciences*, *30*(1), 14–21. doi:10.1016/j.tins.2006.11.005

Pouli, T., Cunningham, D. W., & Reinhard, E. (2011). A Survey of Image Statistics Relevant to Computer Graphics. *Computer Graphics Forum*, *In Press*. doi:10.1111/j.1467-8659.2011.01900.x

Previc, F. H., Declerck, C., & De Brabander, B. (2005). Why your "head is in the clouds" during thinking: the relationship between cognition and upper space. *Acta psychologica*, *118*(1-2), 7–24. doi:10.1016/j.actpsy.2004.10.012

Pridmore, R. W. (2008a). Chromatic induction: Opponent color or complementary color process? *Color Research & Application*, *33*(1), 77–81. doi:10.1002/col.20363

Pridmore, R. W. (2008b). Color constancy from invariant wavelength ratios: I. The empirical spectral mechanism. *Color Research & Application*, *33*(3), 238–249. doi:10.1002/col.20405

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological review*, *117*(4), 1113–43. doi:10.1037/a0020311

Pylyshyn, Z. W. (1999a). Vision and cognition: How do they connect? *Behavioral and Brain Sciences*, *22*(3), 401–414. doi:10.1017/S0140525X99622026

Pylyshyn, Z. W. (1999b). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*(03), 341–65; discussion 366–423. doi:10.1017/S0140525X99002022

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, *80*(1-2), 127–158. doi:10.1016/S0010-0277(00)00156-6

Pylyshyn, Z. W. (2004). Some puzzling findings in multiple object tracking: I. Tracking without keeping track of object identities. *Visual Cognition*, *11*(7), 801–822. doi:10.1080/13506280344000518

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, *3*(3), 179–197. doi:10.1163/156856888X00122

Quaia, C., Optican, L. M., & Goldberg, M. E. (1998). The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Networks*, *11*(7-8), 1229–1240. doi:10.1016/S0893-6080(98)00069-0

Railo, H., & Koivisto, M. (2012). Two means of suppressing visual awareness: a direct comparison of visual masking and transcranial magnetic stimulation. *Cortex*, *48*(3), 333–43. doi:10.1016/j.cortex.2010.12.001

Raizada, R. D. S., & Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: computational clues from the visual system. *Cerebral Cortex*, *13*(1), 100–13. doi:10.1093/cercor/13.1.100

Ramachandran, V. S. (1992a). Blind spots. *Scientific American*, *266*(5), 86–91. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1566041

Ramachandran, V. S. (1992b). Filling in the blind spot. *Nature*, *356*(6365), 115. doi:10.1038/356115a0

Rauschenberger, R. (2003). When something old becomes something new: Spatiotemporal object continuity and attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(3), 600–615. doi:10.1037/0096-1523.29.3.600

Reinhard, E., Adhikhmin, M., Gooch, B., & Shirley, P. (2001). Color transfer between images. *IEEE Computer Graphics and Applications*, *21*(4), 34–41. doi:10.1109/38.946629

Rensink, R. A. (2002). Change detection. *Annual review of psychology*, *53*, 245–77. doi:10.1146/annurev.psych.53.100901.135125

Ridder III, W. H., & Tomlinson, A. (1995). Spectral characteristics of blink suppression in normal observers. *Vision Research*, *35*(18), 2569–2578. doi:10.1016/0042-6989(95)00011-N

Ridder III, W. H., & Tomlinson, A. (1997). A comparison of saccadic and blink suppression in normal observers. *Vision research*, *37*(22), 3171–9. doi:10.1016/S0042-6989(97)00110-7

Rieger, J. W., Grüschow, M., Heinze, H., & Fendrich, R. (2007). The appearance of figures seen through a narrow aperture under free viewing conditions: effects of spontaneous eye motions. *Journal of vision*, *7*(6), 10.1–13. doi:10.1167/7.6.10

Risner, M. L., & Gawne, T. J. (2009). The response dynamics of primate visual cortical neurons to simulated optical blur. *Visual neuroscience*, *26*(4), 411–20. doi:10.1017/S0952523809990174

Ro, T., Breitmeyer, B. G., Burton, P., Singhal, N. S., & Lane, D. (2003). Feedback contributions to visual awareness in human occipital cortex. *Current Biology*, *13*(12), 1038–41. doi:10.1016/S0960-9822(03)00337-3

Roach, N. W., McGraw, P. V, & Johnston, A. (2011). Visual Motion Induces a Forward Prediction of Spatial Pattern. *Current Biology*, *21*(9), 740–745. doi:10.1016/j.cub.2011.03.031

Robinson, A. E., & De Sa, V. R. (2008). Brief presentations reveal the temporal dynamics of brightness induction and White's illusion. *Vision research*, *48*(22), 2370–81. doi:10.1016/j.visres.2008.07.023

Rock, I., Halper, F., DiVita, J., & Wheeler, D. (1987). Eye movement as a cue to figure motion in anorthoscopic perception. *Journal of Experimental Psychology: Human Perception and Performance*, *13*(3), 344–352. doi:10.1037/0096-1523.13.3.344

Roebroeck, A., Formisano, E., & Goebel, R. (2009). The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution. *NeuroImage*, *58*(2), 296–302. doi:10.1016/j.neuroimage.2009.09.036

Roelofs, C. O. (1961). Complementary colour perception. *Acta Psychologica*, *18*(1), 218–228. doi:10.1016/0001-6918(61)90012-9

Rolfs, M., Jonikaitis, D., Deubel, H., & Cavanagh, P. (2011). Predictive remapping of attention across eye movements. *Nature neuroscience*, *14*(2), 252–256. doi:10.1038/nn.2711

Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., Kopell, N. J., et al. (2008). Temporal Interactions between Cortical Rhythms. *Frontiers in neuroscience*, *2*(2), 145–54. doi:10.3389/neuro.01.034.2008

Roorda, A., Metha, A. B., Lennie, P., & Williams, D. R. (2001). Packing arrangement of the three cone classes in primate retina. *Vision research*, *41*(10-11), 1291–306. doi:10.1016/S0042-6989(01)00043-8

Roorda, A., & Williams, D. R. (1999). The arrangement of the three cone classes in the living human eye. *Nature*, *397*(6719), 520–2. doi:10.1038/17383

Rosenbloom, P. S. (2011). Rethinking cognitive architecture via graphical models. *Cognitive Systems Research*, *12*(2), 198–209. doi:10.1016/j.cogsys.2010.07.006

Ross, C. F., & Kirk, E. C. (2007). Evolution of eye size and shape in primates. *Journal of human evolution*, *52*(3), 294–313. doi:10.1016/j.jhevol.2006.09.006

Ross, J., Morrone, M. C., & Burr, D. C. (1997). Compression of visual space before saccades. *Nature*, *386*(6625), 598–601. doi:10.1038/386598a0

Ross, J., Morrone, M. C., Goldberg, M. E., & Burr, D. C. (2001). Changes in visual perception at the time of saccades. *Trends in neurosciences*, *24*(2), 113–21. doi:10.1016/S0166-2236(00)01685-4

Rucci, M. (2008). Fixational eye movements, natural image statistics, and fine spatial vision. *Network (Bristol, England)*, *19*(4), 253–85. doi:10.1080/09548980802520992

Rucci, M., & Beck, J. (2005). Effects of ISI and flash duration on the identification of briefly flashed stimuli. *Spatial vision*, *18*(2), 259–73. doi:10.1163/1568568053320585

Rucci, M., Iovin, R., Poletti, M., & Santini, F. (2007). Miniature eye movements enhance fine spatial detail. *Nature*, *447*(7146), 851–4. doi:10.1038/nature05866

Rudd, M. E., & Arrington, K. F. (2001). Darkness filling-in: a neural model of darkness induction. *Vision research*, *41*(27), 3649–62. doi:10.1016/S0042-6989(01)00216-4

Rumelhart, D. E., McClelland, J. L., & Group, P. R. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations* (p. 567). Cambridge, MA: MIT Press. Retrieved from http://mitpress.mit.edu/books/parallel-distributed-processing

Ryan, J. D., & Villate, C. (2009). Building visual representations: The binding of relative spatial relations across time. *Visual Cognition*, *17*(1), 254–272. doi:10.1080/13506280802336362

Rüter, J., Francis, G., Frehe, P., & Herzog, M. H. (2011). Testing dynamical models of vision. *Vision research*, *51*(3), 343–51. doi:10.1016/j.visres.2010.11.013

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature neuroscience*, *5*(7), 631–2. doi:10.1038/nn876

Sampanes, A. C., Tseng, P., & Bridgeman, B. (2008). The role of gist in scene recognition. *Vision research*, *48*(21), 2275–83. doi:10.1016/j.visres.2008.07.011

Sapkota, R. P., Pardhan, S., & Van der Linde, I. (2011). Object-position binding in visual short-term memory for sequentially presented unfamiliar stimuli. *Perception*, *40*(5), 538–48. doi:10.1068/p6899

Sato, T. K., Nauhaus, I., & Carandini, M. (2012). Traveling Waves in Visual Cortex. *Neuron*, *75*(2), 218–229. doi:10.1016/j.neuron.2012.06.029

Satoh, S., & Usui, S. (2008a). Computational theory and applications of a filling-in process at the blind spot. *Neural Networks*, *21*(9), 1261–71. doi:10.1016/j.neunet.2008.05.001

Satoh, S., & Usui, S. (2008b). Computational Understanding and Modeling of Filling-In Process at the Blind Spot. In M. Ishikawa, K. Doya, H. Miyamoto, & T. Yamakawa (Eds.), *Neural Information Processing* (Vol. 4984, pp. 943–952). Berlin, Heidelberg: Springer. doi:10.1007/978-3-540-69158-7_97

Satoh, S., & Usui, S. (2009). Engineering-approach accelerates computational understanding of V1-V2 neural properties. *Cognitive neurodynamics*, *3*(1), 1–8. doi:10.1007/s11571-008-9065-x

Schill, K., & Zetzsche, C. (1995). A model of visual spatio-temporal memory: the icon revisited. *Psychological research*, *57*(2), 88–102. doi:10.1007/BF00447079

Schindler, A., & Bartels, A. (2013). Parietal Cortex Codes for Egocentric Space beyond the Field of View. *Current biology : CB*, *23*(2), 177–182. doi:10.1016/j.cub.2012.11.060

Schneegans, S., & Schöner, G. (2012). A neural mechanism for coordinate transformation predicts pre-saccadic remapping. *Biological cybernetics*, *106*(2), 89–109. doi:10.1007/s00422-012-0484-8

Scholl, B. J. (2008). What Have We Learned about Attention from Multiple-Object Tracking (and Vice Versa)? In D. Dedrick & L. Trick (Eds.), *Computation, cognition, and Pylyshyn* (pp. 49–78). Cambridge, MA: MIT Press. Retrieved from http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11764

Scholl, B. J., Pylyshyn, Z. W., & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple object tracking. *Cognition*, *80*(1-2), 159–77. doi:10.1016/S0010-0277(00)00157-8

Scholte, H. S., Jolij, J., Fahrenfort, J. J., & Lamme, V. a F. (2008). Feedforward and recurrent processing in scene segmentation: electroencephalography and functional magnetic resonance imaging. *Journal of cognitive neuroscience*, *20*(11), 2097–109. doi:10.1162/jocn.2008.20142

Schreuder, D. (2008). The human observer: visual perception. *Outdoor Lighting: Physics, Vision and Perception* (pp. 273–311). Dordrecht: Springer Netherlands. doi:10.1007/978-1-4020-8602-1_8

Schummers, J., Mariño, J., & Sur, M. (2004). Local networks in visual cortex and their influence on neuronal responses and dynamics. *Journal of physiology, Paris*, *98*(4-6), 429–41. doi:10.1016/j.jphysparis.2005.09.017

Schwartz, E. L. (1983). Cortical mapping and perceptual invariance: a reply to Cavanagh. *Vision research*, *23*(8), 831–5. doi:10.1016/0042-6989(83)90206-7

Schütz, A. C., Braun, D. I., & Gegenfurtner, K. R. (2009). Improved visual sensitivity during smooth pursuit eye movements: temporal and spatial characteristics. *Visual neuroscience*, *26*(3), 329–40. doi:10.1017/S0952523809990083

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301–12. doi:10.1016/j.neuron.2007.06.015

Shea-Brown, E., Gilzenrat, M. S., & Cohen, J. D. (2008). Optimization of decision making in multilayer networks: the role of locus coeruleus. *Neural computation*, *20*(12), 2863–94. doi:10.1162/neco.2008.03-07-487

Sherman, S. M., & Guillery, R. W. (1998). On the actions that one nerve cell can have on another: distinguishing "drivers" from "modulators". *Proceedings of the National Academy of Sciences of the United States of America*, *95*(12), 7121–6. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9618549/

Shiffrin, R. M. (2010). Perspectives on Modeling in Cognitive Science. *Topics in Cognitive Science*, *2*(4), 736–750. doi:10.1111/j.1756-8765.2010.01092.x

Shimojo, S., Wu, D.-A., & Kanai, R. (2003). Coexistence of colour filling-in and filling-out in segregated surfaces. *Perception: ECVP Abstract Supplement*, *32*, 155–155. Retrieved from http://www.perceptionweb.com/abstract.cgi?id=v030892

Shultz, S., Klin, A., & Jones, W. (2011). Inhibition of eye blinking reveals subjective perceptions of stimulus salience. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(52), 21270–5. doi:10.1073/pnas.1109304108

Shuwairi, S. M., Curtis, C. E., & Johnson, S. P. (2007). Neural substrates of dynamic object occlusion. *Journal of cognitive neuroscience*, *19*(8), 1275–85. doi:10.1162/jocn.2007.19.8.1275

Silberstein, M. (2011). Metaphysics or science: The battle for the soul of philosophy of mind. *Philosophical Psychology*, *24*(4), 561–573. doi:10.1080/09515089.2011.583006

Simons, D. J., & Ambinder, M. S. (2005). Change Blindness. Theory and Consequences. *Current Directions in Psychological Science*, *14*(1), 44–48. doi:10.1111/j.0963-7214.2005.00332.x

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception*, *28*(9), 1059–74. doi:10.1068/p2952

Simons, D. J., Lleras, A., Martinez-Conde, S., Slichter, D., Caddigan, E., & Nevarez, G. (2006). Induced visual fading of complex images. *Journal of vision*, *6*(10), 1093–101. doi:10.1167/6.10.9

Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and future. *Trends in cognitive sciences*, *9*(1), 16–20. doi:10.1016/j.tics.2004.11.006

Sligte, I. G., Vandenbroucke, A. R. E., Scholte, H. S., & Lamme, V. a F. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in psychology*, *1*(October), 175. doi:10.3389/fpsyg.2010.00175

Sloman, A. (2009). Some Requirements for Human-Like Robots: Why the Recent Over-Emphasis on Embodiment Has Held Up Progress. In B. Sendhoff, E. Körner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-Like Intelligence* (Vol. 5436, pp. 248–277). Berlin, Heidelberg: Springer. doi:10.1007/978-3-642-00616-6_12

Smith, G., Howell, E. R., & Stanley, G. (1982). Spatial frequency and the detection of temporal discontinuity in superimposed and adjacent gratings. *Perception & Psychophysics*, *31*(3), 293–297. doi:10.3758/BF03202537

Smith, R. A., & Cass, P. F. (1989). Effect of eccentricity on spatial summation and acuity. *Journal of the Optical Society of America A*, *6*(10), 1633. doi:10.1364/JOSAA.6.001633

Snyder, L. H., Calton, J. L., Dickinson, A. R., & Lawrence, B. M. (2002). Eye-hand coordination: saccades are faster when accompanied by a coordinated arm movement. *Journal of neurophysiology*, *87*(5), 2279–86. doi:10.1152/jn.00854.2001

Sohmiya, T., & Sohmiya, K. (1992). Where does an anorthoscopic image appear? *Perceptual and motor skills*, *75*(3 Pt 1), 707–14. doi:10.2466/PMS.75.7.707-714

Sohmiya, T., & Sohmiya, K. (1994). What is a crucial determinant in anorthoscopic perception? *Perceptual and motor skills*, *78*(3 Pt 1), 987–98. doi:10.2466/pms.1994.78.3.987

Sokolov, A., & Pavlova, M. A. (2006). Visual motion detection in hierarchical spatial frames of reference. *Experimental Brain Research*, *174*(3), 477–86. doi:10.1007/s00221-006-0487-6

Solomon, S. G., & Lennie, P. (2005). Chromatic gain controls in visual cortical neurons. *Journal of Neuroscience*, *25*(19), 4779–92. doi:10.1523/JNEUROSCI.5316-04.2005

Sommer, M. A., & Wurtz, R. H. (2008). Visual perception and corollary discharge. *Perception*, *37*(3), 408–418. doi:10.1068/p5873

Song, Y., Makarov, V. A., & Velarde, M. G. (2009). Stability switches, oscillatory multistability, and spatio-temporal patterns of nonlinear oscillations in recurrently delay coupled neural networks. *Biological Cybernetics*, *101*(2), 147–167. doi:10.1007/s00422-009-0326-5

Soo, F. S., Schwartz, G. W., Sadeghi, K., & Berry II, M. J. (2011). Fine spatial information represented in a population of retinal ganglion cells. *Journal of Neuroscience*, *31*(6), 2145–55. doi:10.1523/JNEUROSCI.5129-10.2011

Spering, M., & Carrasco, M. (2012). Similar effects of feature-based attention on motion perception and pursuit eye movements at different levels of awareness. *Journal of Neuroscience*, *32*(22), 7594–601. doi:10.1523/JNEUROSCI.0355-12.2012

Spering, M., Pomplun, M., & Carrasco, M. (2011). Tracking Without Perceiving: A Dissociation Between Eye Movements and Motion Perception. *Psychological Science*, *22*(2), 216–225. doi:10.1177/0956797610394659

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1–29. doi:10.1037/h0093759

Spillmann, L., Otte, T., Hamburger, K., & Magnussen, S. (2006). Perceptual filling-in from the edge of the blind spot. *Vision research*, *46*(25), 4252–7. doi:10.1016/j.visres.2006.08.033

Sporns, O. (2011). The human connectome: a complex network. *Annals of the New York Academy of Sciences*, *1224*, 109–125. doi:10.1111/j.1749-6632.2010.05888.x

Spratling, M. W. (2002). Cortical Region Interactions and the Functional Role of Apical Dendrites. *Behavioral and Cognitive Neuroscience Reviews*, *1*(3), 219–228. doi:10.1177/1534582302001003003

Spratling, M. W. (2004). Neural coding strategies and mechanisms of competition. *Cognitive Systems Research*, *5*(2), 93–117. doi:10.1016/j.cogsys.2003.11.002

St Cyr, G. J., & Fender, D. H. (1969). The interplay of drifts and flicks in binocular fixation. *Vision research*, *9*(2), 245–65. doi:10.1016/0042-6989(69)90004-2

Steinbach, M. J. (1976). Pursuing the perceptual rather than the retinal stimulus. *Vision research*, *16*(12), 1371–6. doi:10.1016/0042-6989(76)90154-1

Steinman, R. M., Haddad, G. M., Skavenski, A. A., & Wyman, D. (1973). Miniature eye movement. *Science*, *181*(102), 810–9. doi:10.1126/science.181.4102.810

Stepp, N., Chemero, A., & Turvey, M. T. (2011). Philosophy for the Rest of Cognitive Science. *Topics in Cognitive Science*, *3*(2), 425–437. doi:10.1111/j.1756-8765.2011.01143.x

Stevenson, S. B., Raghunandan, A., Frazier, J., Poonja, S., & Roorda, A. (2004). Fixation jitter, motion discrimination and retinal imaging. *Journal of Vision*, *4*(11), 85–85. doi:10.1167/4.11.85

Sui, J., Adali, T., Yu, Q., Chen, J., & Calhoun, V. D. (2012). A review of multivariate methods for multimodal fusion of brain imaging data. *Journal of neuroscience methods*, *204*(1), 68–81. doi:10.1016/j.jneumeth.2011.10.031

Surridge, A. K., Osorio, D. C., & Mundy, N. I. (2003). Evolution and selection of trichromatic vision in primates. *Trends in Ecology & Evolution*, *18*(4), 198–205. doi:10.1016/S0169-5347(03)00012-0

Sweet, A. L. (1953). Temporal discrimination by the human eye. *American journal of psychology*, *66*(2), 185–98. doi:10.2307/1418725

Swindale, N. V. (2008). Feedback decoding of spatially structured population activity in cortical maps. *Neural computation*, *20*(1), 176–204. doi:10.1162/neco.2008.20.1.176

Szinte, M., & Cavanagh, P. (2009). Apparent motion from outside the visual field: Retinotopic cortices may register extraretinal locations. *Journal of Vision*, *9*(8), 694–694. doi:10.1167/9.8.694

Szinte, M., & Cavanagh, P. (2012). Apparent motion from outside the visual field, retinotopic cortices may register extra-retinal positions. *PloS one*, *7*(10), e47386. doi:10.1371/journal.pone.0047386

Tadin, D., Lappin, J. S., Blake, R., & Glasser, D. M. (2010). High temporal precision for perceiving event offsets. *Vision research*, *50*(19), 1966–71. doi:10.1016/j.visres.2010.07.005

Tanaka, M. (2007). Spatiotemporal properties of eye position signals in the primate central thalamus. *Cerebral Cortex*, *17*(7), 1504–15. doi:10.1093/cercor/bhl061

Tark, K.-J., & Curtis, C. E. (2009). Persistent neural activity in the human frontal cortex when maintaining space that is off the map. *Nature neuroscience*, *12*(11), 1463–8. doi:10.1038/nn.2406

Tatler, B. W. (2009). Current understanding of eye guidance. *Visual Cognition*, *17*(6), 777–789. doi:10.1080/13506280902869213

Tatler, B. W., & Land, M. F. (2011). Vision and the representation of the surroundings in spatial memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1564), 596–610. doi:10.1098/rstb.2010.0188

Taylor, J. G. (2007). On the neurodynamics of the creation of consciousness. *Cognitive neurodynamics*, *1*(2), 97–118. doi:10.1007/s11571-006-9011-8

Teller, D. Y. (1984). Linking propositions. *Vision research*, *24*(10), 1233–46. doi:10.1016/0042-6989(84)90178-0

Ternus, J. (1926). Experimentelle Untersuchungen über phänomenale Identität. *Psychologische Forschung*, *7*(1), 81–136. doi:10.1007/BF02424350

Thomas, L. E., & Irwin, D. E. (2006). Voluntary eyeblinks disrupt iconic memory. *Perception & psychophysics*, *68*(3), 475–88. doi:10.3758/BF03193691

Thompson, E., & Varela, F. J. (2001). Radical embodiment: neural dynamics and consciousness. *Trends in cognitive sciences*, *5*(10), 418–425. doi:10.1016/S1364-6613(00)01750-2

Thomson, L. C., & Wright, W. D. (1947). The colour sensitivity of the retina within the central fovea of man. *Journal of physiology*, *105*(4), 316–31. Retrieved from www.ncbi.nlm.nih.gov/pubmed/16991729

Thorpe, S. J. (2009). Single units and sensation: still just as relevant today. *Perception*, *38*(6), 802–3; discussion 804–7. doi:10.1068/pmkbar

Thorpe, S. J., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, *381*(6582), 520–2. doi:10.1038/381520a0

Thurgood, C., Whitfield, T. W. A., & Patterson, J. (2011). Towards a visual recognition threshold: New instrument shows humans identify animals with only 1ms of visual exposure. *Vision research*, *51*(17), 1966–1971. doi:10.1016/j.visres.2011.07.008

Tong, J., Stevenson, S. B., & Bedell, H. E. (2008). Signals of eye-muscle proprioception modulate perceived motion smear. *Journal of vision*, *8*(14), 7.1–6. doi:10.1167/8.14.7

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological bulletin*, *215*(3), 216–42. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/19098144

Trautteur, G., Datteri, E., & Santoro, M. (2008). Empirically testable models are needed for understanding visual prediction. *Behavioral and Brain Sciences*, *31*(02), 217–218. doi:10.1017/S0140525X08004019

Treue, S., & Martinez-Trujillo, J. C. (2007). Attending to features inside and outside the spotlight of attention. *Neuron*, *55*(2), 174–6. doi:10.1016/j.neuron.2007.07.005

Troxler, D. (Ignaz P. V. (1804). Über das Verschwinden gegebener Gegenstände innerhalb unseres Gesichtskreises [On the disappearance of given objects from our visual field]. In K. Himly & J. A. Schmidt (Eds.), *Ophthalmologische Bibliothek* (5th ed., Vol. 2, pp. 1–53). Jena, Germany: Fromann. Retrieved from http://books.google.ca/books?id=KRs_AAAAcAAJ

Tse, P. U., Baumgartner, F. J., & Greenlee, M. W. (2010). Event-related functional MRI of cortical activity evoked by microsaccades, small visually-guided saccades, and eyeblinks in human visual cortex. *NeuroImage*, *49*(1), 805–16. doi:10.1016/j.neuroimage.2009.07.052

Tsuchiya, N., & Koch, C. (2009). The Relationship Between Consciousness and Attention. *Neurology of Consciousness*, 63–77. doi:10.1016/B978-0-12-374168-4.00006-X

Tuck, J. P., & Long, G. M. (1990). The role of small-field tritanopia in two measures of colour vision. *Ophthalmic and Physiological Optics*, *10*(2), 195–199. doi:10.1111/j.1475-1313.1990.tb00975.x

Turvey, M. T. (1977). Contrasting orientations to the theory of visual information processing. *Psychological Review*, *84*(1), 67–88. doi:10.1037//0033-295X.84.1.67

Valderrama-gonzalez, G. D., McGinnity, T. M., Maguire, L., & Wu, Q. (2009). Review of Neuron Types in the Retina: Information Models for Neuroengineering. In C. Alippi, M. Polycarpou, C. Panayiotou, & G. Ellinas (Eds.), *Artificial Neural Networks – ICANN 2009* (Vol. 5768, pp. 335–344). Berlin, Heidelberg: Springer. doi:10.1007/978-3-642-04274-4_35

Valdes-Sosa, P. A., Roebroeck, A., Daunizeau, J., & Friston, K. J. (2011). Effective connectivity: Influence, causality and biophysical modeling. *NeuroImage*, *58*(2), 339–361. doi:10.1016/j.neuroimage.2011.03.058

Van de Laar, T., & De Regt, H. C. D. G. (2008). Special Section: Is Cognitive Science Changing its Mind?: Introduction to Embodied Embedded Cognition and Neurophenomenology. *Theory & Psychology*, *18*(3), 291–296. doi:10.1177/0959354308089786

Van den Heuvel, M. P., & Hulshoff Pol, H. E. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity. *European neuropsychopharmacology : the journal of the European College of Neuropsychopharmacology*, *20*(8), 519–534. doi:10.1016/j.euroneuro.2010.03.008

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, *21*(5), 615–28; discussion 629–65. doi:10.1017/S0140525X98001733

Van Roosmalen, P. M. B., Kokaram, A. C., & Biemond, J. (1999). Fast high quality interpolation of missing data in image sequences using a controlled pasting scheme. *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No.99CH36258)* (pp. 3105–3108 vol.6). IEEE. doi:10.1109/ICASSP.1999.757498

Van Zoest, W., Hunt, A. R., & Kingstone, A. (2010). Representations in Visual Cognition: It's About Time. *Current Directions in Psychological Science*, *19*(2), 116–120. doi:10.1177/0963721410363895

Velisavljević, L., & Elder, J. H. (2008). Visual short-term memory of local information in briefly viewed natural scenes: configural and non-configural factors. *Journal of vision*, *8*(16), 8.1–17. doi:10.1167/8.16.8

Vogel, H. (1979). A better way to construct the sunflower head. *Mathematical Biosciences*, *44*(3-4), 179–189. doi:10.1016/0025-5564(79)90080-4

Voges, N., Schüz, A., Aertsen, A., & Rotter, S. (2010). A modeler's view on the spatial structure of intrinsic horizontal connectivity in the neocortex. *Progress in neurobiology*, *92*(3), 277–292. doi:10.1016/j.pneurobio.2010.05.001

Volkmann, F. C. (1986). Human visual suppression. *Vision research*, *26*(9), 1401–16. doi:10.1016/0042-6989(86)90164-1

Volkmann, F. C., Riggs, L. A., Ellicott, A. G., & Moore, R. K. (1982). Measurements of visual suppression during opening, closing and blinking of the eyes. *Vision research*, *22*(8), 991–6. doi:10.1016/0042-6989(82)90035-9

Volkmann, F. C., Riggs, L. A., & Moore, R. K. (1980). Eyeblinks and visual suppression. *Science*, *207*(4433), 900–2. doi:10.1126/science.7355270

Wachtler, T., Doi, E., Lee, T.-W., & Sejnowski, T. J. (2007). Cone selectivity derived from the responses of the retinal cone mosaic to natural scenes. *Journal of vision*, *7*(8), 6.1–14. doi:10.1167/7.8.6

Wade, N. J., & Heller, D. (1997). Scopes of perception: the experimental manipulation of space and time. *Psychological research*, *60*(4), 227–37. doi:10.1007/BF00419407

Wagner, F., Schmuki, R., Wagner, T., & Wolstenholme, P. (2006). *Modeling Software with Finite State Machines: A Practical Approach* (p. 302). CRC Press. Retrieved from http://www.crcpress.com/product/isbn/9780849380860

Wald, G. (1967). Blue-blindness in the normal fovea. *Journal of the Optical Society of America (1917-1983)*, *57*(11), 1289–301. doi:10.1364/JOSA.57.001289

Walker, J. T., & Irion, A. L. (1982). Apparent displacement of moving and stationary strobe flashes. *Human factors*, *24*(2), 213–24. doi:10.1177/001872088202400207

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*(2), 366–83. doi:10.1016/j.neuron.2007.10.012

Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, *60*(2), 215–34. doi:10.1016/j.neuron.2008.09.034

Warren, P. A., & Rushton, S. K. (2009). Optic flow processing for the assessment of object movement during ego movement. *Current Biology*, *19*(18), 1555–60. doi:10.1016/j.cub.2009.07.057

Warren, W. H. (2006). The dynamics of perception and action. *Psychological review*, *113*(2), 358–89. doi:10.1037/0033-295X.113.2.358

Watson, A. B., & Ahumada, A. J. (2011). Blur clarified: a review and synthesis of blur discrimination. *Journal of vision*, *11*(5), 10.1–23. doi:10.1167/11.5.10

Webster, M. A. (2009). Calibrating color vision. *Current Biology*, *19*(4), R150–2. doi:10.1016/j.cub.2008.11.051

Weil, R. S., & Rees, G. (2011). A new taxonomy for perceptual filling-in. *Brain research reviews*, *67*(1-2), 40–55. doi:10.1016/j.brainresrev.2010.10.004

Weisberg, M. (2006). Forty Years of "The Strategy": Levins on Model Building and Idealization. *Biology & Philosophy*, *21*(5), 623–645. doi:10.1007/s10539-006-9051-9

Wertheimer, M. (1912). Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie / Journal of Psychology*, *61*(1), 161–265. Retrieved from http://www.getcited.org/pub/103413947

Westheimer, G. (2007). The ON-OFF dichotomy in visual processing: from receptors to perception. *Progress in retinal and eye research*, *26*(6), 636–48. doi:10.1016/j.preteyeres.2007.07.003

Westheimer, G. (2009). Visual acuity: information theory, retinal image structure and resolution thresholds. *Progress in retinal and eye research*, *28*(3), 178–86. doi:10.1016/j.preteyeres.2009.04.001

Westheimer, G., & Ley, E. J. (1997). Spatial and temporal integration of signals in foveal line orientation. *Journal of neurophysiology*, *77*(5), 2677–84. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9163383

Westheimer, G., & McKee, S. P. (1975). Visual acuity in the presence of retinal-image motion. *Journal of the Optical Society of America (1917-1983)*, *65*(7), 847. doi:10.1364/JOSA.65.000847

Westheimer, G., & McKee, S. P. (1977). Perception of temporal order in adjacent visual stimuli. *Vision Research*, *17*(8), 887–892. doi:10.1016/0042-6989(77)90062-1

White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of vision*, *11*(6), 15.1–10. doi:10.1167/11.6.15

Whitham, E. M., Fitzgibbon, S. P., Lewis, T. W., Pope, K. J., Delosangeles, D., Clark, C. R., Lillie, P., et al. (2011). Visual Experiences during Paralysis. *Frontiers in human neuroscience*, *5*(December), 160. doi:10.3389/fnhum.2011.00160

Whitney, D. (2006). Contribution of bottom-up and top-down motion processes to perceived position. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1380–97. doi:10.1037/0096-1523.32.6.1380

Whitney, D., & Levi, D. M. (2011). Visual crowding: a fundamental limit on conscious perception and object recognition. *Trends in cognitive sciences*, *15*(4), 160–168. doi:10.1016/j.tics.2011.02.005

Wilensky, U., & Resnick, M. (1999). Thinking in Levels: A Dynamic Systems Approach to Making Sense of the World. *Journal of Science Education and Technology*, *8*(1), 3–19. doi:10.1023/A:1009421303064

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of vision*, *4*(12), 1120–35. doi:10:1167/4.12.11

Williams, D. R., MacLeod, D. I. A., & Hayhoe, M. M. (1981). Foveal tritanopia. *Vision Research*, *21*(9), 1341–1356. doi:10.1016/0042-6989(81)90241-8

Willmer, E. N., & Wright, W. D. (1945). Colour Sensitivity of the Fovea Centralis. *Nature*, *156*(3952), 119–121. doi:10.1038/156119a0

Wilson, J. R., & Sherman, S. M. (1976). Receptive-field characteristics of neurons in cat striate cortex: Changes with visual field eccentricity. *Journal of neurophysiology*, *39*(3), 512–33. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/948006

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic bulletin & review*, *9*(4), 625–36. doi:10.3758/BF03196322

Wilson, M. H., & Brocklebank, R. W. (1955). Complementary hues of after-images. *Journal of the Optical Society of America (1917-1983)*, *45*(4), 293–9. doi:10.1364/JOSA.45.000293

Wilson, R. I., & Du Lac, S. (2011). Sensory and motor systems. *Current opinion in neurobiology*, *21*(4), 517–519. doi:10.1016/j.conb.2011.08.001

Winkler, P., McDermott, K. C., Caplovitz, G., & Webster, M. (2011). Figural Chasers. *Journal of Vision*, *11*(11), 1014–1014. doi:10.1167/11.11.1014

Wischnewski, M., Belardinelli, A., Schneider, W. X., & Steil, J. J. (2010). Where to Look Next? Combining Static and Dynamic Proto-objects in a TVA-based Model of Visual Attention. *Cognitive Computation*, *2*(4), 326–343. doi:10.1007/s12559-010-9080-1

Wojtas, D. H., Wu, B., Ahnelt, P. K., Bones, P. J., & Millane, R. P. (2008). Automated analysis of differential interference contrast microscopy images of the foveal cone mosaic. *Journal of the Optical Society of America A*, *25*(5), 1181–9. doi:10.1364/JOSAA.25.001181

Womelsdorf, T., Anton-Erxleben, K., Pieper, F., & Treue, S. (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nature neuroscience*, *9*(9), 1156–60. doi:10.1038/nn1748

Womelsdorf, T., Anton-Erxleben, K., & Treue, S. (2008). Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *Journal of Neuroscience*, *28*(36), 8934–44. doi:10.1523/JNEUROSCI.4030-07.2008

Wu, J.-Y., Xiaoying Huang, & Chuan Zhang. (2008). Propagating waves of activity in the neocortex: what they are, what they do. *Neuroscientist*, *14*(5), 487–502. doi:10.1177/1073858408317066

Wuerger, S. M., Atkinson, P., & Cropper, S. J. (2005). The cone inputs to the unique-hue mechanisms. *Vision research*, *45*(25-26), 3210–23. doi:10.1016/j.visres.2005.06.016

Wunderlich, K., Rangel, A., & O'Doherty, J. P. (2009). Neural computations underlying action-based decision making in the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(40), 17199–204. doi:10.1073/pnas.0901077106

Wurtz, R. H. (2008). Neuronal mechanisms of visual stability. *Vision research*, *48*(20), 2070–89. doi:10.1016/j.visres.2008.03.021

Wyatt, H. J., Pola, J., Fortune, B., & Posner, M. (1994). Smooth pursuit eye movements with imaginary targets defined by extrafoveal cues. *Vision research*, *34*(6), 803–20. doi:10.1016/0042-6989(94)90218-6

Yarbus, A. L. (1967). *Eye Movements and Vision*. (L. A. Riggs, Ed.)*Problems of Information Transmission* (Translated., p. 222). New York: Plenum Press. Retrieved from http://books.google.ca/books?id=9mF2SQAACAAJ

Yellott, J. I. (1983). Spectral consequences of photoreceptor sampling in the rhesus retina. *Science*, *221*(4608), 382–385. doi:10.1126/science.6867716

Yeshurun, Y., & Marom, G. (2008). Transient spatial attention and the perceived duration of brief visual events. *Visual Cognition*, *16*(6), 826–848. doi:10.1080/13506280701588022

Yi, D., Jiang, P., & Mallen, E. (2011). Enhancement of image luminance resolution by imposing random jitter. *Neural Computing & Applications*, *20*(2), 261–272. doi:10.1007/s00521-010-0433-1

Yokoi, I., & Komatsu, H. (2010). Cortico-geniculate feedback linking the visual fields surrounding the blind spot in the cat. *Experimental Brain Research*, *202*(1), 247–51. doi:10.1007/s00221-009-2123-8

Yuste, R. (2011). Review Dendritic Spines and Distributed Circuits. *Neuron*, *71*(5), 772–781. doi:10.1016/j.neuron.2011.07.024

Zeki, S. M. (2005). The Ferrier Lecture 1995 behind the seen: the functional specialization of the brain in space and time. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1458), 1145–83. doi:10.1098/rstb.2005.1666

Zgurovsky, M. Z., & Pankratova, N. D. (2007). Basic Notions of System Analysis. *System Analysis: Theory and Applications* (pp. 35–83). Berlin, Heidelberg: Springer. doi:10.1007/978-3-540-48880-4_2

Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*(4), 423–8. doi:10.1111/j.1467-9280.2009.02322.x

Zhang, W., & Luck, S. J. (2011). The number and quality of representations in working memory. *Psychological science*, *22*(11), 1434–41. doi:10.1177/0956797611417006

Zhou, H., & Desimone, R. (2011). Feature-Based Attention in the Frontal Eye Field and Area V4 during Visual Search. *Neuron*, *70*(6), 1205–17. doi:10.1016/j.neuron.2011.04.032

Zylberberg, A., Dehaene, S., Mindlin, G. B., & Sigman, M. (2009). Neurophysiological bases of exponential sensory decay and top-down memory retrieval: a model. *Frontiers in computational neuroscience*, *3*(March), 4. doi:10.3389/neuro.10.004.2009

Zylberberg, A., Fernández Slezak, D., Roelfsema, P. R., Dehaene, S., & Sigman, M. (2010). The brain's router: a cortical network model of serial processing in the primate brain. *PLoS computational biology*, *6*(4), e1000765. doi:10.1371/journal.pcbi.1000765

Zöllner, F. (1862). Ueber eine neue Art anorthoskopischer Zerrbilder. *Annalen der Physik und Chemie*, *193*(11), 477–484. doi:10.1002/andp.18621931108

Öğmen, H. (2007). A theory of moving form perception: Synergy between masking, perceptual grouping, and motion computation in retinotopic and non-retinotopic representations. *Advances in cognitive psychology / University of Finance and Management in Warsaw*, *3*(1-2), 67–84. doi:10.2478/v10053-008-0015-2

Öğmen, H., & Herzog, M. H. (2010). The Geometry of Visual Perception: Retinotopic and Nonretinotopic Representations in the Human Visual System. *Proceedings of the IEEE*, *98*(3), 479–492. doi:10.1109/JPROC.2009.2039028