

The Embedded Neuron, the Enactive Field?

Mazviita Chirimuuta* & Ian Gold*†

*School of Philosophy & Bioethics, Monash University, Clayton VIC 3800, Australia; mazviita.chirimuuta@arts.monash.edu.au

†Departments of Philosophy & Psychiatry, McGill University, Montreal, Quebec H3A 2T7, Canada; ian.gold@mcgill.ca

Abstract

The concept of the receptive field, first articulated by Hartline, is central to visual neuroscience. The receptive field of a neuron encompasses the spatial and temporal properties of stimuli that activate the neuron, and, as Hubel and Wiesel conceived of it, a neuron's receptive field is static. This makes it possible to build models of neural circuits and to build up more complex receptive fields out of simpler ones. Recent work in visual neurophysiology is providing evidence that the classical receptive field is an inaccurate picture. The receptive field seems to be a dynamic feature of the neuron. In particular, the receptive field of neurons in V1 seems to be dependent on the properties of the stimulus. In this paper, we review the history of the concept of the receptive field and the problematic data. We then consider a number of possible theoretical responses to these data.

1. Introduction

One role for the philosopher of neuroscience is to examine issues raised by the central concepts of neuroscience (Gold & Roskies forthcoming) just as philosophy of biology does for biology and philosophy of physics for physics. In this paper we make an attempt to explore one issue around the concept of the *receptive field* (RF) of visual neurons. Fundamentally, the RF of a neuron represents “how a cell responds when a point of light falls in a position of space (or time)” (Rapela, Mendel & Grzywacz 2006, p. 464). It also describes the kind of stimulus that activates a neuron—a moving bar, a red patch, or whatever. The phrase “receptive field” was coined by the American neurophysiologist and Nobel

laureate, Haldan K. Hartline (1903-1983), in 1938 (Hartline 1938) and has become the central way of characterizing neurons in the visual system and elsewhere. Barlow (1953; see Lettvin, Maturana, McCulloch & Pitts 1968) and, in particular, Hubel and Wiesel (e.g., Hubel & Wiesel 1959) developed the concept. Currently, as Rapela, Mendel, & Grzywacz (2006, p. 464) say, “[r]eceptive fields are the main theoretical framework to represent functional properties of visual cells.”

However, recent findings in the neurophysiology of neurons in primary visual cortex (V1) are at odds with the “classical” conception of the RF (Albright & Stoner 2002), and it is possible that the concept of the visual RF is in transition. Our aim in this paper, therefore, is to examine the concept of the RF and explore the possible consequences for the concept of these data for neurophysiology and computational vision. We are not concerned to argue for a particular view about the status of the concept but rather to begin to articulate some of the options. We make some anodyne remarks at the end of the paper about which of the options we think are most promising, but our purpose here is merely to contribute to the beginning of a discussion about the issues.

We introduce the concept of visual RFs by discussing the classical picture of V1 physiology, most associated with the work of Hubel and Wiesel (section 2). We then turn to the psychophysics and computational vision of contrast discrimination in order to place the visual neurophysiology in context (section 3). We next review the recent data that have raised questions about the classical conception of the RF (section 4). We turn then to consider some of the options available for absorbing the data into visual theory (section 5). We conclude with some remarks on the relevance of these data for thinking about the role of the environment in visual theory (section 6).

2. The Classical Receptive Field

The simplest picture or model of visual processing is of a one-way flow of information from the photoreceptors of the eye, via successively more complex information transformations in the retina, thalamus and the visual cortex towards integration with other functions in the higher cortical areas. Most likely owing to its simplicity, it has long been the picture most appealing to scientists confronted with the daunting complexity of the visual brain. Indeed, this “bottom-up” picture

has held sway even though neuroanatomy has shown, for example, that “top-down” neuronal connections from visual cortex to the thalamus outnumber bottom-up connections from the retina (see Sherman and Guillery 2002). The hierarchical scheme is illustrated in Figure 1a which shows the key anatomical loci of the human visual system up to and including primary visual cortex. These brain structures are shared by higher mammals, and for this reason studies on cat, ferret and other mammals aside from primates, have been considered important steps toward an understanding human vision.

Figures 1b-c illustrate Hubel and Wiesel’s famous hierarchical model of RFs in primary visual cortex, where the elongated RFs of V1 simple and complex cells are taken to be the result of the arrangement of cells lower down in the visual pathway which synapse onto them. The modern era of study of the RF begins with the work of Hubel and Wiesel, and we will, for the purposes of this paper, take theirs to be the standard conception of the RF. However, we note that the work of many other scientists is of equal importance. In particular, Hubel and Wiesel’s studies were of a largely qualitative nature, so almost all quantitative measurement understanding of RF properties is due to the efforts of other laboratories.

Much of the basic understanding of the visual neurons that inspired the traditional picture came from extracellular electrode recordings in the cat. In such experiments, various stimuli would be presented to the eyes of anaesthetised animals while electrodes measured neuronal responses in the cortex in terms of spikes per second. A crucial finding of Hubel & Wiesel (1959) was that neurons in the cat visual cortex respond well to moving or flashing bars of a particular orientation, width and location (see Hubel & Wiesel 1998 for a historical overview of their work, including the “accidental” discovery of elongated RFs). This suggested that in contrast to the neurons of the retina and lateral geniculate nucleus (LGN) which had been found to have circular RFs (Kuffler 1953), V1 RFs were elongated and orientation specific. Hubel and Wiesel mapped these RFs by flashing small spots of light in the visual field. As in earlier studies, any area in space in which flashing a light elicited an increase in neuronal firing rate was defined as part of the ON portion of the RF, and any adjacent area in which a black spot (i.e. a decrease in luminance relative to the background) elicited firing was taken to constitute part of the OFF area. Figure 2a illustrates a receptive field

mapped in this way by Hubel and Wiesel (1959). Hubel & Wiesel (1962) also made a distinction between *simple* and *complex* cells, the first type apparently exhibiting *predictable linear spatial summation* and the second type not. In effect, only simple cell RFs had clearly defined ON and OFF regions and could be mapped with localised spots of light.¹ In contrast, complex cells showed unpredictable *nonlinear* spatial summation and were indifferent to the phase of a bar or grating stimulus (see Figure 3). In the absence of defined ON and OFF regions, the complex cell response would be the same whatever the precise position of the white and dark portions of the stimulus with respect to the RF. These findings motivated the hierarchical model represented in Figure 1b. An obvious explanation for the structure of the simple cell RF is that a small number of LGN cells whose RFs occur in a row in visual space all synapse onto one simple cell. Likewise, an obvious explanation for phase invariance of the complex cells is that these neurons receive input from a small number of simple cells whose RFs overlap but are of different ON-OFF polarity in space.

Other researchers went on to make more detailed studies of simple and complex cells in order to quantify, for example, the linearity of spatial summation. These studies often used sinusoidal grating stimuli rather than spots of light or bars (see Figure 3a). One important reason for using sinusoids was that physiologists were engaged in performing a *systems analysis* of primary visual cortex (see Albrecht, Geisler & Crane 2003). On the assumption that V1 neurons are linear analysers, these methods, borrowed from the physical sciences and engineering, show how the response of such linear neurons to sinusoids can then be used to predict the responses to any image. Recordings were made both in the cat (see, e.g., Henry 1977, Movshon, Thompson & Tolhurst 1978 a,b,c; Jones et al. 1987; Jones & Palmer 1987; Li, Peterson & Freeman 2003) and monkey (e.g., Hubel & Wiesel 1968; Hawken & Parker 1987; Parker & Hawken 1988; Ringach 2002). Figures 2b and c show further examples of simple cell receptive field maps, similar to those presented by Jones, Stepnowski and Palmer (1987) and De Angelis, Ohzawa and Freeman (1993a), respectively. These studies revealed properties of V1 neurons that significantly challenge the classical picture. Indeed, Hubel and Wiesel's original hierarchical model was soon found to be inconsistent with quantitative measurement of simple and complex cell responses implicating some

¹ See Mechler and Ringach (2002) for the case against the simple-complex dichotomy. The classification has faced severe scrutiny but is still in play in the physiology literature. Since this debate is not crucial to our discussion of receptive fields it will not be discussed further here.

role for top-down or lateral connections in shaping RF properties. Discussion of these later studies, and the extent to which they undermine the classical conception of the RF, is the subject of sections 4 and 5. In the next section, however, we examine the influence of the classical conception in the other disciplines of visual neuroscience, notably psychophysics and computation.

3. Psychophysics, Computational Modelling and the Classical RF

3.1. Psychophysics

Psychophysics is a sub-discipline of visual neuroscience in which detailed, quantitative measurements are made of assumed basic visual responses or percepts. For example, classic psychophysical studies measured absolute detection thresholds for dim spots of light, and also for sinusoidal gratings of different spatial frequencies. Historically, much research on the physiology of vision has been motivated by psychophysical findings.² A major subject of psychophysical investigation has been the supposed properties of *spatial frequency* and *orientation channels* (Campbell & Robson 1968; Blakemore & Campbell 1969). A corresponding target for physiologists has been to find a neural explanation for these results (e.g. Campbell, Cooper & Enroth-Cugell 1969; Maffei & Fiorentini 1973). The idea of a channel is basically that of a spatial frequency or orientation selective filter (Braddick, Campbell & Atkinson 1978; Graham 1989), such a filter being the result of the operation of one or more structures in the visual system. Just as any continuous, intensity-varying signal, such as a sound wave, can be described as a set of sinusoidal Fourier components of different amplitudes, frequencies, and phases with respect to one another, any visual image can be analysed as a two-dimensional Fourier transform (Robson 1980; Westheimer 2001). The channel theory of vision was, therefore, the working hypothesis that the visual system itself breaks down images roughly into its Fourier components, by way of its channels, and that, for each channel, there is a sinusoidal stimulus of a particular spatial frequency and orientation to which the channel gives its optimal response (Campbell & Robson 1968).

On Hubel and Wiesel's description of visual cortex, neuronal response properties are fixed, and dependent solely on the response properties of the upstream

² Though this is more true of the British than the American school of visual neuroscience (Lennie & Movshon 2005)

neurons which provide their input. A number of authors (e.g. Blakemore & Campbell 1969, Campbell et al. 1969, Maffei & Fiorentini 1973) were therefore prompted to equate these cortical properties with the properties of the psychophysical channels. However, as Marr & Hildreth (1980) pointed out, these physiological and psychophysical theories of cortical processing are rather different in that the psychophysical channels are said to perform something akin to a Fourier transform of the visual image which is a *non-local* analysis of frequency; the simple cells of Hubel and Wiesel, on the other hand, operate as detectors of *localised* contrast features, such as edges. Still, the channel hypothesis is now established, in so far as it is generally accepted that the key mechanisms in the visual system revealed by psychophysics are spatial frequency and orientation selective, rather than broadband (Majaj, Pelli, Kurshan & Palomares 2002).

At the same time, channel models have evolved. Originally, it was not thought that the response of one channel should alter the output of another channel (but see Tolhurst 1972). But in response to more recent neurophysiological work (see section 4.1) and in order to better account for psychophysical data, some psychophysicists have rejected the *independent*-channels hypothesis, developing models in which channels are dynamically effected by the responses of other channels (e.g. Foley 1994). The convergence of psychophysics and neurophysiology has also been given a helping hand in recent years with the advancement of scanning techniques. In particular, with functional magnetic resonance imaging (fMRI) it is possible for experimenters to track areas of increased neural activity whilst observers perform traditional psychophysical tasks. Boynton, Demb, Glover & Heeger (1999) argue that observers' contrast discrimination functions (detection of an increase in grating contrast as a function of background contrast) can be predicted by fMRI signals in V1 and V2, implicating these areas as critical for setting thresholds in this task. This result gave new support to psychophysicists' attempts to account for their data in terms of V1 physiology (Foley 1994, Chirimuuta and Tolhurst 2005).

3.2. Computational Models of V1 RFs

As noted in section 2, Hubel and Wiesel's work was largely qualitative, but scientists following them aimed to get a mathematically precise grasp of V1 physiology. A powerful tool here was the computational modelling of RFs. A

computational model of an RF is supposed to capture the key functional properties of the RF, such that it can be used to predict how a neuron will respond to any hypothetical stimulus. With the intense research on V1 following Hubel and Wiesel, there soon grew to be a large body of data on response properties, and some of these data sets gave conflicting evidence on key questions such as whether or not there really are two classes of simple and complex cells. One way of usefully integrating this large amount of data was to develop computational models of cortical receptive fields and to measure the goodness of fit with the data: if the fit is good, it may be inferred that the mathematical principle operational in the model (e.g. linear versus nonlinear contrast response) captures the key properties of the neuron.

A large number of simple cell models have been developed since the 1980s. One which has achieved notable popularity is the Gabor model. The Gabor function is the product of a sinusoid and a Gaussian envelope, giving a localised sinusoidal modulation. The one-dimensional Gabor function was first developed by Dennis Gabor (1946) for use in communications engineering, and it is a particularly useful coding function because it minimises joint uncertainty about time and temporal frequency (Gabor-Heisenberg-Weyl uncertainty). In the case of visual analysis, the function minimises joint uncertainty about location and spatial frequency, enabling one to perform *local* Fourier analysis (see section 3.1 above). The function was introduced to vision science by Marcelja (1980).

With the Gabor model comes the implication that simple cells are essentially linear fixed filters whose job it is to analyse any given visual scene into simple bar or blob like components. Thus the Gabor model shares a common fate with a particular conception of the RF. Over the past two decades the appropriateness of the Gabor as a model for V1 neurons has been researched intensively and, given the significant nonlinearities reported, it seems a fair summary of the findings to say that the Gabor model can account for roughly half of the response behaviour of simple cells in anaesthetized animals (DJ Tolhurst, personal communication; see Carandini et al 2005 for a recent assessment of the standard RF models). In section 4 we will discuss some of these reported nonlinearities, and in section 5 we will discuss neuroscientists' responses to the reported discrepancies between the data and the linear model.

3.3. Computer Vision

The field of computer vision, as opposed to the development of computational techniques in neurophysiology, attempts to reproduce useful visual function in computers or robots. Perhaps the most influential figure in computer vision is David Marr. His project was to develop mechanisms that were equivalent to the biological ones, and that could operate in artificial systems (Marr 1982). Crucial to his methodology was the distinction between *algorithm* and *implementation*, or “software” and “hardware”. This distinction allowed him to argue that a process in the visual system, for example, making a selective response to vertical edges, could be exactly equivalent to a computational process such as convolution with vertical filters, even though the processes are realised in very different physical substrates. Marr hoped to find algorithms which could carry out processes useful to machine vision, such as edge extraction for the purpose of object recognition (Marr & Hildreth 1980). Still, what is shared by computer vision and the computational modelling discussed in section 3.2 is the assumption that the processing that takes place in the visual system can also be implemented in a digital computer. David Marr’s work was not especially inspired by detailed physiology, but later researchers have taken this line further. For example, John Daugman’s algorithm for iris scanning is a convergence of ideas from computer vision, statistics, V1 physiology and computational modelling of V1 (see Daugman 2003).

4. Recent Findings in V1 Physiology

Hubel and Wiesel’s conception of the RF is known as the “classical RF” because further investigation of V1 physiology revealed that a given neuron’s response could be modified by stimulation of the neuron in visual field regions which would not in themselves elicit a response, or by presentation of stimuli to which the neuron was apparently unresponsive. Such findings challenged the picture of the visual system as feedforward and hierarchical, with little or no modulation of responses due to interaction between neurons at the same level of the hierarchy, or from higher levels. Another assumption of the classical picture was that RFs are fixed properties of neurons. This has also been challenged by recent work. This section reviews some key findings in the extensive literature on the visual cortex.

In passing, it is worth asking to what extent these discoveries have been made

possible with new techniques unavailable in the 1950s and 1960s. For example, the advent of *intracellular* recording allowed researchers to record directly inhibitory input to V1 neurons, now taken to be a critical factor behind RF tuning properties. It has also been argued that recordings from awake-behaving animals have revealed nonlinearities not apparent in the traditional anaesthetised preparation (Lamme, 2004). On the other hand, it is worth considering the idea that complex nonlinear behaviour is receiving more attention now because scientists' conception of V1 function and RFs has altered, making more salient complex behaviours which might previously have been put down to noise in the system (see section 4.4).

4.1. Inhibitory Networks and Surround Effects in V1

Later neurophysiological investigation did not bear out the conjecture of Hubel & Wiesel's (1962) hierarchical model according to which response properties of cortical neurons can be explained in terms of summation of upstream neurons which have simpler RFs. For example, the intracellular recordings of Hirsch, Alonso, Reid & Martinez (1998) found that cortical neurons receive a significant amount of inhibitory and excitatory input from within the cortex, as well as the excitatory geniculate input mentioned by the hierarchical model. Furthermore, computational studies (Troyer, Krukowski, Priebe & Miller 1998, Lauritzen, Krukowski & Miller 2001, Wielaard, Shelley, McLaughlin & Shapley 2001) have shown that the inhibitory input is necessary for keeping tuning bandwidth invariant with stimulus contrast, as is approximately the case in V1 (Sclar & Freeman 1982, Skottun, Bradley, Sclar, Ohzawa & Freeman 1987).

It follows from these findings that cortical response properties cannot be independent of the activity of neighbouring neurons (Blakemore & Tobin 1972). An important example is the work of Bonds (1989), which showed that simultaneous stimulation with an optimal stimulus and a superimposed mask at a different orientation or spatial frequency, will cause the neuron's responses to drop below its response level to the optimal stimulus alone. Since the neuron is not thought to be directly (i.e. by way of excitatory geniculate input) affected by the mask to which it is poorly tuned, the implication is that the neuron is receiving inhibition from neurons which *do* respond to the mask.

Another line of research which has challenged the "discrete" receptive fields

picture is the investigation of the effects of stimulating beyond the spatial extent of the “classical receptive field”. Hubel and Wiesel (1962) defined the receptive field as the area over which a neuron responds to small spots of light.³ However, it has been shown that V1 neurons will often produce a greater response if a stimulus is extended beyond this area, even if stimulation in this area alone is not able to drive the neuron. The maximum extent of the area which causes progressive excitation is known as the “summation field”. Stimulation beyond the “summation field” often causes a decline in response and the area over which one observes this inhibition is known as the “suppressive surround”. Examples of different surround effects can be found in the work of, amongst others, Blakemore & Tobin (1972), Maffei & Fiorentini (1976), Nelson & Frost (1978), Gilbert & Wiesel (1990), De Angelis, Freeman & Ohzawa (1994), Jones, Grieve, Wang & Sillito (2001), Cavanaugh, Bair & Movshon (2002a), Cavanaugh, Bair & Movshon (2002b) and Levitt & Lund (2002); see Albright & Stoner (2002) and Tucker & Fitzpatrick (2003) for reviews.

The existence of a suppressive surround means that neurons are affected by parts of the image adjacent to their receptive fields and so, in an ecological context where there will be complex image structure around the receptive field (rather than blank grey screen), it will be difficult to predict the responses to any particular stimulus. There has been much speculation over the purpose of the surround in ecological vision. Marcus & van Essen (2002) suggest that the surround may aid scene segmentation in primate V1 and V2; similarly, Li & Gilbert (2002) and Sugita (1999) suggest a role in contour integration and grouping problems (see Lamme (2004) for a review). Following such results, another new concept that has been added to that of the RF, is that of the “association field” (Kapadia, Westheimer & Gilbert 2000). It was a term first introduced in psychophysics (Field, Hayes & Hess 1993), but in neurophysiology the association field maps the amount of modulation that is invoked by different stimuli surrounding an optimal stimulus for any given RF.

³ C.f. Barlow et al (1967). This, however, is just one way of measuring the size of the classical RF; it measures the “minimum response field” (MRF). Another method is to stimulate the neuron with an grating of increasing size to find the optimal stimulus dimensions for its RF, measuring the “grating summation field” (GSF) (DeAngelis et al. 1994; Sceniak et al. 1999). The first method tends to give smaller estimates of RF size than the other; see Cavanaugh et al. 2002a.

To conclude this subsection, we note that the finding of the interdependence of neurons' RF properties raises questions about what level of analysis – single neuron or population – is best for experimental work in the visual cortex, an issue we will raise again in section 5.

4.2. The Dynamic RF

As noted above, the traditional concept of a RF is of a fixed filter – that is, of a unit that signals the presence of its preferred stimulus, its preference unaffected by recent history of stimulation or by the activity of other units. This notion is tied up with the idea that visual neurons represent the features to which they are responsive, or that they perform pattern recognition (Craik 1966). The physiological findings that we will review in this section cast doubt on the assumption that RF properties are fixed. As Tucker and Fitzpatrick (2003) have recently put it, “[t]he cortical RF has become a dynamic entity, one in which context and history play significant roles in shaping its boundaries and altering its properties.”

The question we will address in section 5 is how revisionary or conservative neuroscientists could now be about the concept of RF in the light of these findings, and in section 6 we will ask if the dynamic V1 neuron points to new ways of thinking about perception in general. But we begin a discussion of this literature with one of the most robust reports on how RF size is readily modified by stimulus contrast. A number of different research groups at around the same time all reported that the CRF is found to be larger if the neuron is stimulated with low contrast gratings (Levitt & Lund 1997, Polat, Mizobe, Pettet, Kasamatsu & Norcia 1998, Kapadia, Westheimer and Gilbert 1999 and Sceniak, Ringach, Hawken & Shapley 1999). Explanations for this phenomenon suggest that it is a means of increasing sensitivity at low contrasts, analogous to the way in which photoreceptors and other cells in the retina show increased pooling of signals between neighbouring neurons in dim light conditions. The conceptual interest of this result is simply that it means one cannot speak of an RF as having a fixed size; size must always be given relative to the contrast at which the neurons was tested, and this complicates the traditional model of the RF.

A comparable result to this is the finding that RF size is also dependent on the neuron's recent history of stimulation. “Artificial scotoma” is the term given to a

blank stimulus presented in the RF centre which suppresses response activity. Gilbert and Wiesel (1992), Kapadia, Gilbert and Westheimer (1994) and Pettet and Gilbert (1992) have shown that presentation of an artificial scotoma for a number of minutes causes the RF of cortical neurons to grow to several times their original size. The effect is reversible by subsequent presentation of stimuli to which the neuron is responsive and is thought to be mediated by horizontal connections between neurons in the same cortical layer (Tucker and Fitzpatrick 2003). Indeed, fast plasticity of horizontal connections, and also of top-down connections from higher cortical areas, are commonly put forward as the physiological explanation of RF dynamism, and of the surround and cross stimulus effects discussed above, though the details of such mechanisms remain an area of contention.

One striking demonstration of the dynamism of visual neuronal properties can be seen in the work of Bair and Movshon (2004) on direction-selective (DS) neurons in V1 and motion area MT/V5 of the macaque monkey. Such neurons have been modelled extensively as linear filters which are oriented in spacetime to give directional sensitivity, and in such models RFs are taken to be stable. However, these authors note many psychophysical reports showing that at a perceptual level the *temporal profile of motion integration* – the time course or pattern of motion analysis undertaken by the visual system – is variable with stimulus speed, spatial frequency and contrast. So the aim of the investigation was to find out if this variability is a property which arises at a population level, while the profiles of individual neurons are fixed, or if the basis of the variability can be shown at the level of individual neurons whose RFs are dependent on stimulus properties.

Bair and Movshon showed convincingly that the latter is the case. For example, neurons extend their integration time (i.e. the time window in which a spiking response signifies the presence of a stimulus moving in the preferred direction) for slowly moving gratings. They term this “adaptive temporal integration”, since such a stimulus dependent shift is advantageous, improving the signal-to-noise ratio of the response to slowly moving objects. They conclude that “[i]t is possible that no single RF profile can be attributed to a cortical cell. This implies that models relying on a fixed filter to endow component neurons with their tuning properties could be highly inaccurate, in general.”

4.3. *Natural Images*

All of the studies mentioned above have used artificial stimuli, but since the 1990s studies using natural stimuli – photographs or video clips taken in the outside world – have become increasingly central to vision science (see Figure 3b). The reason for this interest in natural stimuli is that the visual system evolved in the natural environment, and presumably many of its features are adaptations to the peculiarities of natural visual information. Of particular interest is the view that properties of simple cell receptive fields are special adaptations to the informational “redundancies” of natural images,⁴ in that they minimise both the correlations between neurons’ responses and level of activity of individual neurons; this is known as *sparse coding* (Baddeley & Hancock 1991, Olshausen & Field 1997, van Hateren & van der Schaaf 1998, Vinje & Gallant 2000, Willmore & Tolhurst 2001).

A key question is whether or not the cortex shows radically different physiological properties under natural and artificial stimulation. Two studies (Ringach, Hawken & Shapley 2002, Smyth, Willmore, Baker, Thompson & Tolhurst 2003) have shown that receptive field maps generated from responses to natural images resemble the elongated, oriented fields derived from the classic grating experiments. In contrast, David, Vinje and Gallant (2004) have made the case that RF models generated by stimulation with natural stimuli give significantly better predictions of responses to novel natural stimuli than do RF models generated by stimulation with gratings. Likewise, the predictions of responses to novel grating stimuli were superior if the RF model had been constructed from the correlations of responses to gratings. If V1 neurons were linear filters, this would not be the case; RF models would generalise between classes of stimuli. So the stimulus specificity of RF predictions that David et al (2004) report, which seem to be due to crucial nonlinearities in V1, again points to a more dynamic notion of RF than was originally conceived.

4.4. *Science and Simplicity*

⁴ Attneave (1954) and Barlow (1960) introduced the idea of redundancy reduction as a design principle of sensory systems. Redundancy reduction may explain why there are so many nonlinearities in V1 responses: “Linear operations can only partially exploit the statistical redundancies of natural scenes, and nonlinear operations are ubiquitous in visual cortex. However, neither the detailed function of the nonlinearities nor the higher-order image statistics are yet fully understood” (Zetzsche & Nuding 2005).

The introduction to this section of the paper raised the question of whether the increase in attention physiologists now give to the complex nonlinear properties of RFs can be put down to the advent of new techniques to reveal such properties, or if a changing conception of the RF has made such properties now more salient to scientists. The answer is probably both. Anecdotally, it is worth noting that Kuffler (1953) presented a messier picture of the responses of visual neurons than either Hartline before him or Hubel and Wiesel after him. In fact, his description of the physiology of the ganglion cells in the cat retina sounds more like the picture emerging from the recent results that we have discussed in this section. For example, he notes that these neurons' response patterns vary with overall illumination changes; that "[t]he most outstanding feature in the present analysis is the flexibility and fluidity of the discharge patterns arising in each RF" (p. 61); and that "[t]here seems to exist a very great variability between individual RFs and therefore a detailed classification cannot be made at present" (p. 62).

Thus, there may not be anything so new after all in the idea of the dynamic RF that has been presented in this section as the result of novel findings. Aside from data made accessible by new technologies, such observations were available to physiologists in the early days of visual neuroscience. This is not to condemn the scientists who put forward the simpler, so-called traditional picture of the RF, for it should be appreciated that the most promising route for any new science has always been to seek out any underlying simplicity in what appears to be a formidably complex and unpredictable object of investigation. Indeed, one wonders if work on visual cortex would have expanded and flourished in the way that it did had Hubel and Wiesel not presented such an attractively neat picture of its physiology. The challenge raised by the recent work discussed in this section is whether these simplifying assumptions ultimately defeated the aim of understanding V1 function by disregarding as noise the very neuronal properties that make our visual system work in the real world as it does.

5. Options

In this section we consider a number of ways the stimulus-dependence data could be integrated into visual neuroscience and cognitive theory. Our aim is to map out some theoretical strategies rather than to defend one in particular. It is premature to be defending one theoretical option when considerably more empirical work is necessary in order to confirm that the classical conception of the RF is in fact

untenable. The job for the philosopher of neuroscience — at least at this stage — is to consider the pros and cons of various ideas as a way of beginning the debate. We consider six options below, in order of increasing radicalness.

5.1. First Option: Expanding the Classical RF

The data reviewed above show that the classical conception of the RF is no longer compelling, but that is not yet to say that the concept of the RF is dead. If the concept can be altered or expanded to accommodate the new data, then this might be the most appropriate strategy to adopt. Since the concept of the RF has proven so useful up to now, better to stretch the concept than to dispense with it. The central issue to be decided is whether the data would require us to stretch the concept to the point where it would no longer be recognizable. Can we revise the concept of the RF or must we eliminate it?

Elegant and conservative extensions of the classical RF may be available. An example is David Heeger's normalisation model of visual neurons. Heeger (1992) notes that the linear model of V1 fails to account for all of the physiological data. Rather than rejecting the linear model outright, Heeger's model includes "divisive normalisation." This is the idea that every neuron's response is divided by a term reflecting the combined activity of all of its neighbours. Thus the local activity is accounted for in a model of the single neuron by means of an equation which summarizes the effects of the circuit, without specifically parametising other neurons. The model does not include biological detail, though it has been noted that the division could be implemented in the brain by what is known as *shunting inhibition* (Carandini, Heeger & Movshon 1997; but see Carandini, Heeger & Senn 2002, Freeman, Durand, Kiper & Carandini 2002, Meier & Carandini 2002). The normalisation model preserves a linear conception of the RF and successfully predicts a good deal of neurophysiological data, including the shape of the contrast response curve, especially response saturation (Maffei & Fiorentini 1973, Albrecht & Hamilton 1982, Sclar et al. 1990), cross-orientation masking (Bonds 1989) and surround suppression (reviewed in Fitzpatrick 2000). Heeger's approach has been particularly influential in subsequent psychophysical models, an important example being Foley's (1994) model of contrast discrimination.

Despite the successes of Heeger's normalisation model in accounting for a number of surround inhibition effects, it by no means assimilates all of

problematic findings listed above. In particular, the normalisation model cannot explain findings of surround *enhancement* (Maffei & Fiorentini 1976, Gilbert & Wiesel 1990), an effect equivalent to the presence of a “summation field” beyond the “classical receptive field.

5.2. Second Option: Piecemeal Solutions

The option just discussed does not tackle head on the problem of stimulus dependence. It may be that a sophisticated extension of the classical RF model will eventually be able to predict that an individual neuron will demonstrate different response properties under natural as opposed to artificial stimulation or as a result of different surround input. On the other hand, no such “universal model” may be forthcoming. Even if the RF changes with categories of stimulus and with visual tasks, it is nonetheless possible to produce models for particular stimulus-task pairs, or classes of pairs. One could give up on the idea of there being a single model of V1 suitable for all stimuli and visual tasks and focus attention on piecemeal solutions to visual problems. This option would allow one to preserve the procedure for modelling visual circuits by treating neurons as fixed (i.e., not stimulus-dependent) components, but the modelling would now be relativized to stimulus class, or visual task, or both.

This solution might be criticised as inelegant and ad hoc; but piecemeal solutions are, in a sense, the norm in visual neuroscience, where RF models are usually devised to account for specific data sets (a particular animal responding to a particular sort of stimulus). It is usually hoped that the model will generalise to novel data, but expectations are that the fit will worsen.

A deeper worry is that by relativising RFs to stimulus classes or visual tasks, we are ignoring some of the significant dynamic features of neurons – precisely those features that are responsive to stimulus or visual task. It would be counter-productive at best to retain the classical receptive field at the expense of ignoring the neuron’s dynamic properties.

5.3. Third Option: Natural Stimuli

Another possible answer to the problem of the stimulus dependency of RFs might be to choose a canonical stimulus. In this case, natural stimuli are the prime

candidate for two reasons: first, because our ultimate goal is to understand vision in the real world; and, second, there is evidence that richer neural responses are revealed by natural stimuli. The prospect of refiguring all of visual neuroscience using natural scenes as canonical stimuli, however, is daunting to say the least because it would require that much of the work of the last forty years be repeated with natural stimuli instead of sinusoids and the like.

Moreover, the idea that RF properties can only be fully revealed by natural stimuli is a radical one. It is sometimes suggested (e.g. David et al 2004) that the RFs revealed by natural stimuli are so complex that a natural image derived model is only good at predicting responses to other natural stimuli. Adopting natural images as canonical stimuli thus entails giving up on the idea that once one fully characterises a neuron's RF, the characterization can then be used to predict responses to *any* stimuli. As Rust and Movshon (2005) put it, “[u]ltimately, one hopes to integrate all these models into a single theory that can predict neuronal and population responses to any arbitrary stimulus.” Even the piecemeal option optimistically leaves open the possibility that all of the partially successful models might be integrated into a powerful general model. To choose natural scenes as canonical stimuli is to give up on this hope.

5.4. Fourth Option: The Primacy of Circuitry

The last two options were presented as only fairly radical. Yet there is a case to be made that the notion of generality is so crucial to the concept of the RF that to give up on it is to change the concept beyond recognition. If we are poised to do mortal damage to the idea of the RF, a natural question to consider is whether we could dispense with the concept of the RF altogether. One way of doing this would be to attempt to model the circuitry of V1 as a whole and hope that the information contained in the RF will emerge or be replaced by something equally informative in the circuit model. The success of Hubel and Weisel's model of V1 was one of the decisive factors that moved vision science in the direction of single-unit, rather than network, analysis (Churchland & Sejnowski 1992). Perhaps a return to the network level is the way forward.

In particular, if we model V1 at the level of its circuitry, then it is possible that the uniformity that would be lost by relativizing RFs to stimulus and task would be recovered. This would be the case if the RF of a particular neuron was altered due

to predictable responses of other V1 neurons in response to the stimulus or task. In other words, if the dynamic nature of the RF could be explained by static features of other V1 neurons and the connections among them, then it is plausible that a single model of V1 could be developed that would have dynamic RFs as a consequence. Perhaps this is what Bair (2005, p. 463) has in mind when he predicts that

The primacy of the RF as a concept for embodying the function of V1 neurons will be replaced by a set of circuits and synaptic mechanisms as our computational models begin to explain ever more response properties. The RF can then be understood as an emergent property that changes with the statistics of the input.

Whether this is the case, of course, is an empirical question. There is no guarantee that V1 circuitry will not itself be affected by feedback from other visual areas and, as a result, prove as dynamic as the RFs of V1 neurons.

Even if we could produce a model of this kind, however, it is worth considering whether it might not be preferable to continue to analyse visual function at the level of individual neurons whose RFs are dynamic rather than to move up to the level of circuitry. Barlow (1972) adapted the phrase “neuron doctrine” (originally the name for the view that the brain is composed of discrete cells rather than being a single continuous structure) to express the view that brain function is best understood at the level of individual neural activity. And one motivation for retaining the notion of a dynamic RF might be to hold on to the neuron doctrine as a general methodological principle of neuroscience. If we reject the notion of a dynamic RF in favour of network- or circuit-level explanation of function, then we must abandon the neuron doctrine in favour of a “higher” level of explanation. Although this might turn out to be a beneficial break from the past, it will require a dramatic rethinking of neural computation. We turn to this issue in the next section.

5.5. Fifth Option: Decoupling Computer Vision from Neurophysiology

The last option considered amounts to an elimination of the concept of the RF, so the question arises of what visual neuroscience would look like without it. We don't know the answer to this question, but it seems certain that it would require a

significant shift in neurophysiological theory and would have consequences for computational approaches to vision. Analysis by fixed filters is a very natural way to think about vision, and it is straightforward to implement artificially. Without the simplifying assumptions behind traditional V1 physiology, the problem of vision begins to look intractable.

In addition, even if one could model a dynamic RF, it is a further question whether it could easily be incorporated into robot vision. We have claimed that at least in some branches of the discipline computational vision and visual neurophysiology have been developed with an eye to linking the two (Teller, 1984) or understanding how visual computations are implemented by neurophysiological mechanisms. One virtue of the classical RF is that it is a natural way to begin to see neurons as elements in these implementations. As we have seen, modelling of early visual processes and the neurophysiology of V1 both conceive of vision as composed of a circuit composed of a small number of simple feed-forward mechanisms, and the history of the study of contrast perception provides ample evidence of the co-evolution of the two disciplines. However, the loss of a fixed RF might lead to a break between the theories of neuronal and artificial vision. One might, therefore, be concerned by the consequences for computational vision of radically altering or giving up the classical conception of the RF. If neurophysiology does not provide computational modelling with the basic units needed, then it is more difficult to model early vision, or, at any rate, to model it with confidence in producing realistic models.

With this kind of worry in mind, one theoretical option would be to give up the idea of an isomorphism between the elements of computational modelling and neurophysiology. On this view, the task of computational modeling remains to provide a theoretical framework for early vision, and the task of visual neurophysiology remains to describe the properties of the neurons and circuits that implement this computation. However, on this view, it is no longer an assumption that the building blocks of the implementation are individual neurons. Although this may be a significant departure from current practice, the idea that individual neurons implement the basic computational processes is a methodological assumption rather than a substantiated doctrine – though it is an assumption which is supported by evidence and is not arbitrary. If we give up the

neuron doctrine, however, then we need not assume that the abstract organization of visual computation has to be implemented by individual neurons that behave as the computational components do. Linear feed-forward models of visual computation need not be implemented by linear feed-forward visual neurons.

5.6. *Which Option?*

Which of the above options should we choose? That question is in part an empirical one but only in part. Some of these paths are likely to be less fruitful than others, so it is worth thinking about which one should choose prior to investing a lot of time in any of them. We suspect that more than one option ought to be pursued and that others are worth setting to one side for the time being. In this section we make some remarks about how to choose among the options.

The first option — that of redescribing or expanding the concept of the RF so as to retain a large part of the classical conception — is also clearly worth pursuing. As we noted, the work of Heeger (1992) shows that there are mathematical techniques that might make this possible. Heeger's work cannot handle all the data, and it is an empirical question whether a single model will do the job, but this is clearly an option that ought to be explored on the grounds of conservatism.

As we remarked with respect to the second option of piecemeal solutions to visual modelling, we think this is an option that is not only inelegant but would leave quite a bit of visual functioning unexplained. We want to know *why* neurons respond in different ways to different stimuli. If we know that, then it's unlikely that piecemeal solutions will remain piecemeal. If we understand how visual neurons change their state and why, then the models of function that are specific to different classes of stimuli will presumably form part of a single unified theory. Although piecemeal exploration is an indispensable dimension of scientific practice, it should not, on our view, be a methodological ideal.

It would be hard to deny that more work with natural stimuli is a crucial requirement of future experiments, and there is already a lot of work being done along these lines. One important question for this work is whether there are a number of clearly distinguishable classes of natural stimuli that produce categorically different neural responses and, if so, how these classes ought to be characterized. We suggested above that forty years of visual neurophysiology

might have to be repeated, but that is a worst-case scenario. It is possible that the use of natural stimuli would produce data that would lead to different ways of conceiving of V1 neurons – and other visual neurons– and generate rapid progress. Nonetheless, it seems quite likely that working with natural stimuli will still require that we think about how to understand the RF in a new way.

The fourth option of exploring circuitry is likely to be a fruitful one. We suspect that this option has not been more fully explored because it is both technically and mathematically complex. It is also possible that a bias in favour of single unit explanations has also influenced the course of research. Dealing with the complexity of the problem is an empirical one, and the encouragement of philosophy is neither necessary nor particularly helpful. In contrast, we think that philosophers *can* argue for calling the neuron doctrine into question and making room for the possibility of circuit-level explanations. Whether such explanations will be successful remains to be seen, but we should not allow the success of single unit neurophysiology to discourage visual neuroscience from thinking of the RF as a derivative (or, as Bair puts it, as an “emergent”) property.

We are also in favour of disengaging computation from neurophysiology to some extent. It is a truism that one of essential features of a computational theory of any kind is that it is “implementable” by neurons. This constraint has become more important in modelling over the years in part (we suspect) as a result of a backlash against the sort of computational theory that saw the brain as an afterthought — as a matter of “mere” implementation. This is a positive development. However, at this stage, we know so little about how the visual system implements visual computation that we should not let this constraint exert too *much* force on computational modelling. By trying to maintain an isomorphism between computation and neurophysiology, computational theory is restricted. In turn, this restriction reduces the ways of thinking about how visual neurons might implement the computation. Different styles of computational modeling might suggest ways of thinking of circuits as the unit of implementation and this might lead to different ways of doing the neurophysiology.

This suggests that the option of changing computational vision to mirror the new data coming from neurophysiology may be one best left to one side for the present as well. Visual modelling might get some ideas from neurophysiology, but we

should not make an isomorphism between the two *de rigueur*. When we know more at both the physiological and modelling levels, the question of implementation can more fruitfully be pursued.

6. Perception and Environment

We noted above that we do not yet know what the mechanism is that underlies stimulus dependence of V1 cells. One possibility is that changes in neural state or responsiveness is a kind of adaptation. It is well known that neurons change their behaviour when the visual stimulus changes substantially. When one moves from conditions of low light levels to bright sunshine, for example, a relatively quick process of adaptation to the light level occurs in which the activity of visual neurons is dampened. The same sort of thing happens with colour. If you put on rose-colored glasses, the world looks pink for a while but soon resumes its chromatic range. Or again if you look at a waterfall, and then look at a stationary scene, the stationary scene seems to move as an after-effect of the adaptation of motion-sensitive neurons to the downward movement of the water. And so on.

Adaptation is important in perception because by altering the state of the visual system, or some part of it, the system is able to function across a greater range of stimuli than would otherwise be possible. A change in neural state or responsiveness to the *statistics* of the visual stimulus might represent a complex version of adaptation (Grzywacz & Balboa 2002) that would expand the repertoire of the visual system and allow it to function effectively in different kinds of natural environments. If something like this is correct, then stimulus-dependence, though surprising, may not be a qualitatively new phenomenon. It does, however, allow us to think somewhat differently about the role of the environment in modulating perception.

There is a longstanding tension between two traditions in the philosophy of perception. One tradition, favoured by analytic philosophy, takes mental representations as central to perception. The perceptual action is all in the mind of the perceiver. In the other tradition, favoured by continental philosophy, perception is more a matter of action than of representation, and, for this reason, the environment in which perception occurs is essential to understanding perception. Merleau-Ponty (1962) is perhaps the most important figure associated with this view.

In recent years, there has been a rapprochement between these traditions. Gareth Evans's solution to the Molyneux problem (Evans, 1996) makes use of something much like Merleau-Ponty's framework for perception. And more recent work, such as that of Clark (1998), has emphasized the importance of the environment in understanding perception. More importantly, neuroscience may be in the process of resolving the debate. The work of Milner and Goodale (2006) has provided evidence that there are two distinct, but interacting, visual systems, one responsible for representing the way the world looks and the other responsible for providing information about how to interact with the objects visually perceived. As Ennen (2003) notes in a different context, the difference between analytic and continental philosophy of perception may be a difference in subject matter (i.e. which of the two visual systems one is of greatest interest) and not in theory.

A recent attempt to reconcile the two traditions in the context of colour perception is due to Thompson, Palacios & Varela (1992). They emphasize the importance of environmental features in developing a theory of colour perception, and they develop an ontological theory of colour that takes colour to be a relation between perceiver and environment. Summarizing Levins and Lewontin (1985) they say:

(1) Organisms determine in and through their interactions what in the physical environment constitutes their relative environments; (2) organisms alter the world external to them as they interact with it; (3) organisms transduce the physical signals that reach them, and so the significance of these signals depends on the structure of the organism; (4) organisms transform the statistical pattern of environmental variation in the world external to them; and (5) the organism-environment relationship defines the "traits" selected for in evolution (cf. Oyama 1985). (p. 21)

They go on to say:

We must encompass both the extradermal world conceived as the animal's environment and the sensory-motor structure of the animal in any adequate theory of perception. (p. 22)

This view, though attractive, is rather programmatic. The notion of stimulus-

dependence as adaptation provides us with one concrete way of thinking about one narrow aspect of visual perception and its relation to the environment. If the mechanism of stimulus-dependence is a kind of adaptation, then the environment can modulate the visual system in a quite complex way by means of a familiar type of mechanism. The statistics of the visual stimulus can alter the way the visual system processes the information it is receiving, and this shows that the statistical properties of the environment are an ineliminable part of a theory of visual perception. V1 neurons may thus give us the beginnings of a theory of the complex interactions between perceiver and environment.

References

- Albrecht D G & Hamilton D B 1982. Striate cortex of monkey and cat: contrast response function. *Journal of Neurophysiology*, 48;217–237.
- Albrecht, D. G., Geisler, W.S. & Crane, A.M. (2003) Nonlinear properties of visual cortex neurons: Temporal dynamics, stimulus selectivity, neural performance. In: L. Chalupa and J. Werner (Eds.), *The Visual Neurosciences*. Boston: MIT Press, 747-764.
- Albright, T.D. & Stoner, G.R. 2002. Contextual influences on visual processing. *Annual Review of Neuroscience* 25:339:379.
- Attneave, F. 1954. Some informational aspects of visual perception, *Psychological Review*, 61, 183-93.
- Baddeley R J & Hancock P J 1991. A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proceedings of the Royal Society of London, B*, 246:219–223
- Bair, W. & Movshon, J.A. 2004. Adaptive Temporal Integration of Motion in Direction-Selective Neurons in Macaque Visual Cortex. *Journal of Neuroscience* 24:7305-7323.
- Bair, W. 2005. Visual receptive field organization. *Current Opinion in Neurobiology* 15:459-464.
- Barlow, H.B. 1953. Summation and inhibition in the frog's retina. *Journal of Physiology* 119: 69-88.
- Barlow, H.B. 1960. The coding of sensory messages, in *Current Problems in Animal Behaviour*, (ed.) W.H. Thorpe and O.L. Zangwill, Cambridge: Cambridge University Press.
- Barlow, H.B. 1972. Single units and sensation: A neuron doctrine for perceptual psychology. *Perception* 1:371-394.
- Barlow HB, Blakemore C, & Pettigrew JD. 1967 The neural mechanism of

binocular depth discrimination. *Journal of Physiology* 193: 327–342

Blakemore, C. & Campbell, F.W. 1969. On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology* 203:237-260.

Blakemore C & Tobin E A 1972. Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research* 15:439–440.

Bonds A B 1989. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience* 2:41–55.

Boynton, G.M., Demb, J.B., Glover, G.H. & Heeger, D.J. 1999. Neuronal basis of contrast discrimination. *Vision Research* 39:257-269.

Braddick, O.J., Campbell, F.W. & Atkinson, J. 1978. Channels in vision: basic aspects. In *Handbook of sensory physiology, vol. 8 (Perception)*, R. Held, H. Leibowitz & H. Teuber (eds) pp.1-38. Heidelberg: Springer.

Campbell, F.W. & Robson, J.G. 1968. Application of fourier analysis to the visibility of gratings *Journal of Physiology* 197:551-566.

Campbell, F.W. Cooper, G.F. & Enroth-Cugell, C. 1969. The spatial selectivity of the visual cells of the cat *Journal of Physiology* 203:223-225.

Carandini M, Heeger D J & Movshon J A 1997. Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, 17:8621–8644.

Carandini M, Heeger D J & Senn W 2002. A synaptic explanation of suppression in visual cortex. *Journal of Neuroscience* 22:10053–10065.

Carandini, M., Demb, J.B., Mante, V., Tolhurst, D.J. Dan, Y., Olshausen, B.A., Gallant, J.L. and Rust, N.C. 2005. Do We Know What the Early Visual System Does? *The Journal of Neuroscience* 25:10577–10597

Cavanaugh J R, Bair W & Movshon J A 2002a. Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *Journal of*

Neurophysiology 88:2530–2546.

Cavanaugh J R, Bair W & Movshon J A 2002b. Selectivity and spatial distribution of signals from the receptive field surround in macaque v1 neurons. *Journal of Neurophysiology* 88:2547–2556.

Chirimuuta, M. & Tolhurst, D.J. 2005. Does a Bayesian model of V1 contrast coding offer a neurophysiological account of human contrast discrimination? *Vision Research* 45:2943-2959

Churchland, P. S. and Sejnowski, T.1992. *The computational brain*. Cambridge, MA, MIT Press.

Clark, A. 1998. *Being There*. Cambridge: MIT Press.

Craik, K.J.W. 1966. *The Nature of Psychology*. Cambridge: Cambridge University Press.

Daugman J 2003 The importance of being random: Statistical principles of iris recognition. *Pattern Recognition* 36:279-291.

S. V. David, W. E. Vinje, and J. L. Gallant 2004 Natural Stimulus Statistics Alter the Receptive Field Structure of V1 Neurons. *Journal of Neuroscience* 24:6991-7006

De Angelis, G.C., Ohzawa, I. and Freeman, R.D. 1993. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. general characteristics and postnatal development. *Journal of Neurophysiology* 69:1091-1117

De Angelis G C, Freeman R D & Ohzawa I 1994. Length and width tuning in the cat's primary visual cortex. *Journal of Neurophysiology* 71:347–374.

Ennen, E. 2003. Phenomenological coping skills and the striatal memory system. *Phenomenology and the Cognitive Sciences* 2:299-325.

Evans, G. 1996. Molyneux's question. In Evans, *Collected Papers*. New York:

Oxford University Press.

Field DJ, Hayes A, & Hess R. 1993 Contour Integration by the Human Visual System: Evidence for a Local 'Association Field'. *Vision Research* 33:173-193.

Fitzpatrick D 2000. Seeing beyond the receptive field in primary visual cortex. *Current Opinion in Neurobiology* 10:438–443.

Foley, J.M. 1994. Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America, A* 11:1710-1719.

Freeman T C B, Durand S, Kiper D C & Carandini M 2002. Suppression without inhibition in visual cortex. *Neuron*, 35:759–771.

Gabor, D. 1946. Theory of communication. *Journal of the Institution of Electrical Engineers* 93:429-459.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A.B. & Massey, J.T. 1989. Mental rotation of the neuronal population vector. *Science* 243:234-236.

Gilbert C D & Wiesel T N 1990. The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, 30:1689–1701.

Gilbert CD & Wiesel TN. 1992 Receptive field dynamics in adult primary visual cortex. *Nature*. 356:150–152

Gold, I. & Roskies, A. forthcoming. Philosophy of neuroscience. In Ruse, M. (ed.) *Oxford Handbook of Philosophy of Biology*.

Graham, N. 1989. *Visual Pattern Analyzers*. Oxford: Clarendon Press.

Balboa, R. M., & Grzywacz, N. M. 2002 A Bayesian framework for sensory adaptation. *Neural Computation* 14:543-559.

Hartline, H.K. 1938. The response of single optic nerve fibers of the vertebrate

- eye to illumination of the retina. *American Journal of Physiology* 121:400-415.
- Hawken, M.J. & Parker, A.J. 1987. Spatial properties of neurons in the monkey striate cortex. *Proceedings of the Royal Society of London, B* 231:251-288
- Heeger D J 1992. Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9:181-197.
- Henry, G.H. 1977. Receptive field classes of cells in the striate cortex of the cat. *Brain Research* 133:1-28.
- Hirsch, J.A., Alonso, J., Reid, R.C. & Martinez, L.M. 1998. Synaptic integration in striate cortical simple cells. *Journal of Neuroscience* 18:9517-9528.
- Hubel, D.H. & Wiesel, T.N. 1959. Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology* 148:574-591.
- Hubel, D.H. & Wiesel, T.N. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* 160:106-154.
- Hubel, D.H. & Wiesel, T.N. 1968. Receptive fields and functional architecture of monkey striate cortex.. *Journal of Physiology* 195:215-244.
- Hubel, D.H. & Wiesel, T.N. 1998. Early exploration of the visual cortex. *Neuron* 20:401-412.
- Jones H E, Grieve K L, Wang W & Sillito A M 2001. Surround suppression in primate V1. *Journal of Neurophysiology* 86:2011-2028.
- Jones, J.P. & Palmer, L.A. 1987. An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology* 58:1233-1258
- Jones, J.P., Stepnowski, A. & Palmer, L.A. 1987. The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology* 58:1212-1232

Kapadia, MK, Gilbert, CD and Westheimer, G. 1994 A quantitative measure for short-term cortical plasticity in human vision. *Journal of Neuroscience* 14: 451-457

Kapadia, MK, Westheimer, G and Gilbert, CD. 1999 Dynamics of spatial summation in primary visual cortex of alert monkeys. *Proceedings of the National Academy of Science USA*, 96: 12073-12078

Kapadia, MK, Westheimer, G and Gilbert, CD. 2000 The spatial distribution of contextual interactions in primary visual cortex and in human perception. *Journal of Neurophysiology* 84: 2048-2062

Kuffler, S.W. 1953. Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology* 16:37-68.

Lamme, V.A.F. 2004. Separate neural definitions of visual consciousness and visual attention. *Neural Networks* 17: 861-872

Lauritzen, J.S., Krukowski, A.E., & Miller, K.D. 2001. Local correlation-based circuitry can account for responses to multi-grating stimuli in a model of cat V1. *Journal of Neurophysiology* 86:1803-1815.

Lennie, P. & Movshon, J. A. 2005. Coding of color and form in the geniculostriate visual pathway. *Journal of the Optical Society of America A* 22:2013-2033

Lettvin, J.Y., H. R. Maturana, H.R., McCulloch, W.S. & Pitts, W.H. 1968. What the frog's eye tells the frog's brain. In Corning, W.C. & Balaban, M. (eds.) *The Mind: Biological Approaches to its Functions*, New York: Interscience Publishers, pp. 233-258.

Levins, R. & Lewontin, R. 1985. *The Dialectical Biologist*. Cambridge MA: Harvard University Press.

Levitt J B & Lund J S 1997. Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73-76.

- Levitt J B & Lund J 2002. The spatial extent over which neurons in macaque striate cortex pool visual signals. *Visual Neuroscience* 19:439–452.
- Li W & Gilbert C D 2002. Global contour saliency and local colinear interactions. *Journal of Neurophysiology* 88:2864–2856.
- Li, B., Peterson, M.R. & Freeman, R.D. 2003. Oblique effect: a neural basis in the visual cortex. *Journal of Neurophysiology* 90:204-217.
- Maffei, L. & Fiorentini, A.. 1973. The visual cortex as a spatial-frequency analyzer. *Vision Research* 13:1255-1267.
- Maffei L & Fiorentini A 1976. The unresponsive regions of visual cortical receptive fields. *Vision Research* 16:1131–1139.
- Majaj, N.J., Pelli, D.G., Kurshan, P. & Palomares, M. 2002. The role of spatial channels in letter identification. *Vision Research* 42:1165-1184.
- Marcelja, S. 1980. Mathematical description of the responses of simple cortical cells. *Journal of the Optical Society of America* 70:1297-1300.
- Marcus D S & van Essen D C 2002. Scene segmentation and attention in primate cortical areas V1 and V2. *Journal of Neurophysiology* 88:2648–2658.
- Marr, D. 1982. *Vision*. San Francisco: W.H. Freeman and co.
- Marr, D. & Hildreth, E. 1980. Theory of edge detection. *Proceedings of the Royal Society of London, B* 207:187-217.
- Mechler, F. & Ringach, D.L. 2002. On the classification of simple and complex cells. *Vision Research* 42:1017-1033.
- Merleau-Ponty, M. 1962. *Phenomenology of Perception*. London: Routledge & Kegan Paul.
- Meier L & Carandini M 2002. Masking by fast gratings. *Journal of Vision*, 2:293–

301.

Milner, D. & Goodale, M. 2006. *The Visual Brain in Action*, 2nd edition. New York: Oxford University Press.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. 1978a. Spatial summation in the receptive fields of simple cells in the cat's striate cortex.. *Journal of Physiology* 283:53-77.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. 1978b. Receptive field organization of complex cells in the cat's striate cortex. *Journal of Physiology* 283:79-99.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. 1978c. Spatial and temporal contrast sensitivity of neurons in areas 17 and 18 of the cat's visual cortex. *Journal of Physiology* 283:101-120.

Nelson J I & Frost B 1978. Orientation selective inhibition from beyond the classical receptive field. *Brain Research* 139:359–365.

Olshausen B A & Field D J 1997. Sparse coding with an overcomplete basis set: a strategy employed by v1? *Vision Research*, 37:3311–3325.

Oyama, S. 1985. *The ontogeny of information*. New York: Cambridge University Press.

Parker, A.J. & Hawken, M.J. 1987. Two-dimensional spatial structure of receptive fields in monkey striate cortex. *Journal of the Optical Society of America, A* 6:598-605

Pettet. M.W. & Gilbert, C.D. 1992 Dynamic Changes in Receptive-Field Size in Cat Primary Visual Cortex *Proceedings of the National Academy of Sciences*, 89:8366-8370

- Polat U, Mizobe K, Pettet M W, Kasamatsu T & Norcia A M 1998. Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature*, 391:580–584.
- Rapela, J., Mendel, J.M., & Grzywacz, N.M. 2006. Estimating nonlinear receptive fields from natural images. *Journal of Vision* 6:441–474, <http://journalofvision.org/6/4/11/>.
- Ringach, D.L. 2002. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology* 88:455-463.
- Ringach D L, Hawken M J & Shapley R 2002. Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *Journal of Vision* 2:12–24.
- Robson, J.G. 1980. Neural images: the physiological basis of spatial vision. In *Visual Coding and Adaptability*, C. Harris (ed.) pp.177-214. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Sceniak M P, Ringach D L, Hawken M J & Shapley R M 1999. Contrast's effect on spatial summation by macaque v1 neurons. *Nature Neuroscience* 2:733–739.
- Sclar G & Freeman R 1982. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research* 46:457–461.
- Sclar G, Maunsell J H R & Lennie P 1990. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, 30:1–10.
- Shepard, R.N. & Cooper, L.A. 1982. *Mental Images and Their Transformations*. Cambridge: MIT Press.
- Sherman SM & Guillery RW 2002. The role of the thalamus in the flow of information to the cortex. *Philosophical Transactions of the Royal Society, Biological Sciences*, 357:1643–1894.
- Skottun B C, Bradley A, Sclar G, Ohzawa I & Freeman R D 1987. The effects of

contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behaviour. *Journal of Neurophysiology* 57:773–786.

Smyth D, Willmore B, Baker G E, Thompson I D & Tolhurst D J 2003. The receptive field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *Journal of Neuroscience*, 23:4746–4759.

Sugita Y 1999. Grouping of image fragments in primary visual cortex. *Nature* 401:269–272.

Teller D Y 1984. Linking propositions. *Vision Research*, 24:1233–1246.

Thompson, E , Palacios, A & Varela, FJ 1992 Ways of Coloring: Comparative Color Vision as a Case Study for Cognitive Science *Behavioral and Brain Sciences* 15: 1–26.

Tolhurst D J 1972. On the possible existence of edge detector neurons in the human visual system. *Vision Research*, 12:797–804.

Troyer, T.W., Krukowski, A.E., Priebe, N. & Miller, K.D. 1998. Contrast-invariant orientation tuning in cat visual cortex: Thalamocortical input tuning and correlation-based intracortical connectivity. *Journal of Neuroscience* 18:5908–5927.

Tucker, T.R. and Fitzpatrick, D. 2003. Contributions of vertical and horizontal circuits to the response properties of neurons in primary visual cortex. In *The Visual Neurosciences*, Chalupa, L.M. and Werner, J.S., eds., MIT Press.

van Hateren J H & van der Schaaf A 1998. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London, B* 256:359–366.

VinjeWE & Gallant J L 2000. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276.

Westheimer G. 2001. The fourier theory of vision. *Perception* 30:531–541

Wielaard, D.J., Shelley, M., McLaughlin, D. & Shapley, R.M. 2001. How simple cells are made in a nonlinear network model of the visual cortex. *Journal of*

Neuroscience 21:5203-5211.

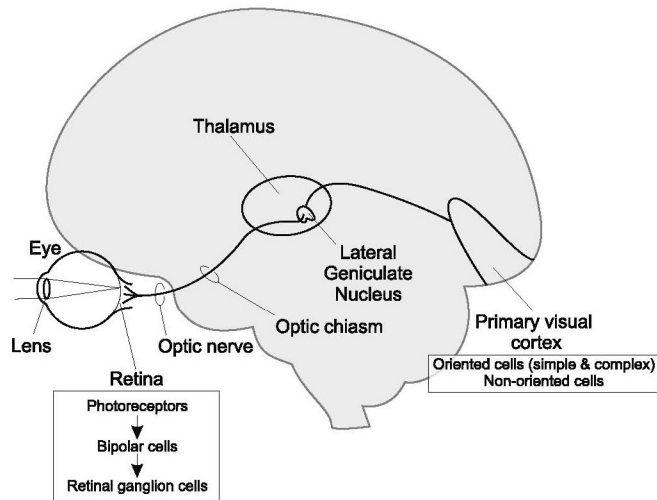
Willmore B & Tolhurst D J 2001. Characterizing the sparseness of neural codes.
Network 12:255–270.

Zetsche & Nuding 2005. *Network: Computation in Neural Systems*
June/September 2005; 16(2/3): 191–221.


Figures

Figure 1

(a)



(b)



(c)




Figure 2




Figure 3

(a)



(b)



Figure Legends

Figure 1

- (a) The main structures of the early mammalian visual system
- (b) From Hubel and Wiesel (1962), the explanation of simple cell elongated receptive fields in terms of the rectangular arrangement of LGN input cells which have circular RFs (left).
- (c) From Hubel and Wiesel (1962). Complex cells were classified as having responses indifferent to the phase (black or white polarity) of the stimulus. This property was explained by their having simple cell inputs with various phase tunings (left).

Figure 2

Three different sorts of receptive field maps of a V1 simple cell. (a) is a qualitative map of the type used by Hubel & Wiesel (1959). Triangles and crosses represent ON and OFF regions, respectively. (b) is a quantitative map, as used by Jones et al. (1987). The height of the surface at each point is proportional to the strength of the cell's response to stimulation at that point, with positive values indicating ON responses and negative values indicating OFF responses. (c) is a quantitative map of the type used by De Angelis et al. (1993a). The brightness at each point is proportional to the strength of the cell's response at that point; mid-grey indicates zero response, brighter shades indicate ON responses and darker shades indicate OFF responses.

Figure 3

Examples of images used as stimuli in neurophysiological and psychophysical experiments.

- (a) A standard sinusoidal grating. In neurophysiological work gratings are normally presented "drifting" across the phase of the sinusoid, rather than static
- (b) Natural image.