

In the interest of saving time: a critique of discrete perception

Tomer Fekete^{*,†}, Sander Van de Cruys, Vebjørn Ekroll and Cees van Leeuwen

Brain and Cognition Unit, KU Leuven, Tiensestraat 102, Leuven, 3000, Belgium

*Correspondence address. Brain and Cognition Unit, KU Leuven, Leuven 3000, Belgium. E-mail: tomer.fekete@gmail.com

†Tomer Fekete, <http://orcid.org/0000-0003-3164-125X>

Abstract

A recently proposed model of sensory processing suggests that perceptual experience is updated in discrete steps. We show that the data advanced to support discrete perception are in fact compatible with a continuous account of perception. Physiological and psychophysical constraints, moreover, as well as our awake-primate imaging data, imply that human neuronal networks cannot support discrete updates of perceptual content at the maximal update rates consistent with phenomenology. A more comprehensive approach to understanding the physiology of perception (and experience at large) is therefore called for, and we briefly outline our take on the problem.

Key words: perception; contents of consciousness; theories and models; psychophysics

Introduction

Recently, Michael Herzog *et al.* (2016) proposed a two-stage model of sensory processing, in which perceptual experience is updated in discrete steps. To be sure, their treatise is thought-provoking, as it offers useful considerations for the study of sensory experience—namely that behavioral methods cannot inform us directly of the actual duration of percepts. Yet, from their arguments it does not follow that percepts are, in a relevant sense, discrete.

We want to challenge their line of thought on several fronts: first, the data professed to support discrete perception do, in fact, turn out to be compatible with continuous-time models of perception. Second, even if perception “were” discrete, this could only marginally differentiate from continuity, given the multiplexing typically occurring in experience, and the physiological and psychophysical conditions under which the perceptual process evolves. We then show that the neural mechanism envisaged for implementing discrete percepts—convergence into attracting steady states—isn’t quite up to the task. We conclude by briefly suggesting that a more comprehensive approach to understanding the physiological underpinning of perception (and experience at large) is called for, and briefly outline our take on the problem.

While Herzog *et al.*’s article is the primary addressee, the scope of our critique extends to a broader range of discrete perception arguments, including, e.g. VanRullen and Koch (2003), VanRullen (2016), and Freeman (2006), as everyday situations call for update rates of perceptual content that our neuronal networks cannot plausibly support to be discrete. In order for this manuscript to be self-contained, we begin by briefly outlining the original proposal.

One Step at a Time

According to Herzog and coworkers, stimuli are analyzed quasi-continuously and unconsciously for features such as orientation or color, but also in order to assign temporal labels such as duration. To integrate these features, typically long periods of unconscious processing are needed before a conscious percept could result. Vision is an ill-posed problem that can only be solved through the time-consuming processes by which, e.g. recurrent neuronal networks settle into steady states. When unconscious integration is finished (a process that may take up to 400 ms, Scharnowski *et al.* 2009), all features are rendered conscious at once. Consciousness, they suggest, will thus remain

Received: 22 May 2017; Revised: 31 January 2018. Accepted: 5 February 2018

© The Author(s) 2018. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

Highlights

- Discrete perception necessitates neural networks to switch states at rates consistent with experience.
- Linear analysis suggests that local networks cannot support switching at plausible rates.
- The asynchronous nature of inter area communication imposed by delays exacerbates the problem.
- Hierarchical, multiscale computation implemented through continuous dynamics near criticality is an alternative.
- Such models do not suffer from the aforementioned drawbacks and naturally accommodate our experience as we know it.

unchanged until the next update or, alternatively, exist for just a single moment until the next update. Updates do not occur at regular intervals, but rather whenever the system settles into an attracting steady state. The time this takes depends on the complexity of the problem. We are not consciously aware of these discrete updates, simply because by definition this would go beyond the temporal resolution of the system.

During the time window until the next update, time is effectively frozen; therefore, we are unable to sense duration as such, but the ongoing unconscious processes may accumulate timekeeping information, allowing an event to be assigned a duration label. As this indirect labeling constitutes our experience of duration, the duration of experience cannot be studied experimentally directly. The authors illustrate their point with the example of the feature fusion paradigm combined with TMS (Scharnowski et al. 2009): A TMS pulse administered “after” stimulus presentation at any point within a window of ~400 ms can change the resulting percept, depending on the exact moment of stimulation. Consequently, they claim, 400 ms is the lower bound for the minimal duration of percepts under this experimental paradigm.

How Postdictive is Too Postdictive

To provide a clear-cut piece of evidence that perception is discrete, Herzog et al. offer the color-phi phenomenon. This phenomenon occurs when a circle is briefly presented first in one location, then in another, meanwhile changing color, e.g. from green to red (Fig. 1). The authors argue: “logically it is impossible to experience the color change before having seen the second disk. The conscious percept must have been formed retrospectively, thus contradicting continuous theories” “... visual awareness is postdictive and thus seemingly incompatible with continuous theories.”

We have no qualms with the postdictive character of the percept—as the perception of any stimulus will necessarily be postdictive given that signal transduction, transmission and analysis require time. Moreover, objects in the natural environment afford signals to different modalities that propagate towards our sensory apparatus at varying speeds (e.g. sound and light). Thus, to make sense of our surroundings, a constant reconciling of information streams is called for. This indeed requires a window of integration that is postdictive (Van Wassenhove et al. 2007; Changizi et al. 2008). There is no reason to think, however, that the occurrence of such a window must imply discreteness. A “sliding window” (i.e. one that is continually updated) would be perfectly suitable for such integration processes. Thus, postdictive integration is perfectly consistent with continuous perception, at least on the face of it.

Indeed, in a series of studies, Eagleman, Sejnowski and colleagues showed for several visual illusions, such as the flash-lag and line motion illusions, that sensory events following the offset of a stimulus can affect its perception (Eagleman and Sejnowski

2000; Eagleman and Sejnowski 2003; Stetson et al. 2006). In a complementary modeling work by Rao et al. (2001), they showed that such results are best explained by a smoothing model. In such a model, current decisions about visual features are informed both by the past (filtering) and immediate future (smoothing) of a sensory event (i.e. sensory activation). The model they present could be implemented either continuously, or by sufficiently dense discrete approximation (~50 Hz in their implementation). They note that the delay of the smoothing buffer is likely a time dependent variable that is modulated by context (e.g. task demands), and is on the order of 100 ms.

Be that as it may, let us examine whether the framework suggested by Herzog et al. indeed explains color-phi better than a continuous framework. The authors argue that the color-phi illusion requires a temporal window of integration spanning at least one stimulus cycle (350 ms, Kolers and von Grünau 1976), after which a percept coding for the entire sequence is rendered conscious at once. However, if a single state encodes for a process, this comes at the price of either (i) claiming that we do not actually experience movement—i.e. experiencing the moving object as being at distinct locations at distinct times—or (ii) admitting that there is perceptual change (e.g. in the location of the object) not mirrored in neuronal change. The latter would amount to violation of supervenience—the notion that consciousness is determined by physical processes—and therefore, that consciousness is amendable to scientific explanation. Accordingly, a discrete rendition of the ensuing percept that does justice to the experience of a moving disc changing its color approximately half way through would require update of perceptual content at several points along the way (see Fig. 1 and Supplementary Figs S1–S3),¹ and not once in bulk at each cycle as Herzog et al. are suggesting. By definition, a continuous framework does not run into these problems.

Does discreteness nevertheless offer any added value in explaining the color-phi phenomenon? It should be noted that apparent motion is just as postdictive as the color change during the illusion: its experience too depends on the presentation of the second stimulus (Nadasdy and Shimojo 2010). What is surprising in color-phi is that color unlike position or shape (as reported in the original study, Kolers and von Grünau 1976) is not interpolated, but shifts abruptly. Discreteness however does not offer any insight as to this curiosity.

1 This is not to be construed as claiming that “time needs to be represented by time” (Dennett and Kinsbourne 1992). Indeed, we agree with Herzog et al. that valuation of passage of time is a distinct feature of our experience in itself, and therefore, at least to some extent, separable from the experience of movement (change). However, experience qua being realized by neuronal processes is associated with a matter of fact regarding events experienced before, together with, or after one another (which is not to say that our “assessment” of our experience must be accurate). This ordering relationship must necessarily hold for the neuronal events realizing the associated phenomenology.

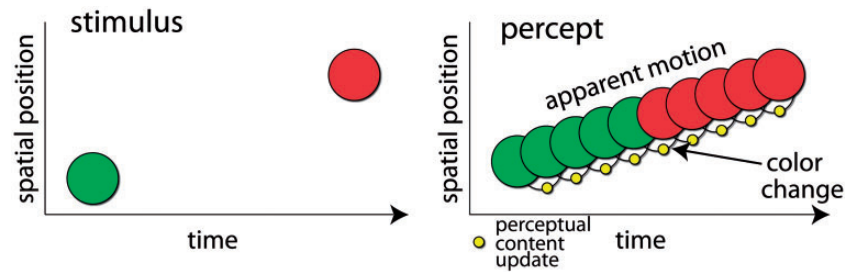


Figure 1. The color phi illusion. Left: the stimulus sequence during one cycle of the illusion. Right: the corresponding percept. The percept comprises a moving colored circle switching color mid-way. From a discrete perspective this requires that the observer go through k distinct perceptual events: i.e. the circle is at position x_i at time t_i with the color C_i for $i = 1, 2, \dots, k$. This would entail $k + 2$ perceptual content updates (blank screen to first circle, circle i to $i + 1$ and so on). That would mean that while experiencing each cycle of the illusion, the corresponding brain activity must change to bring about the rapidly updating perceptual content, by switching through $k + 1$ discrete states. Even though switching might be irregular in time, any such k implies a (n approximate) cortical switching rate. In the example shown $k = 10$. Thus, for a stimulus cycle in this paradigm this would imply at least a rate of 29 Hz of discrete state switching (and see [Supplementary Fig. S1](#)). To substantiate prima facie plausibility for a discrete theory, it is necessary to specify plausible upper bounds for k that would square well with our experience (e.g. while watching a movie or a flying bird, listening to music or speech ...). Further still, it is necessary to show that cortical networks are indeed capable of implementing this maximal rate of state switching.

A continuous outlook, besides squaring better with experience, has no difficulty in explaining interpolation, whether or not it is delayed relative to the stimulus. In fact there is direct imaging data suggesting that apparent motion might result from lateral spread of activity: when V1 is imaged during presentation of stimuli leading to the line motion illusion, similar cascades of activation result from presenting the “real” and “illusory” stimuli ([Jancke et al. 2004](#)). It is therefore plausible to assume that this will lead to similar cascades of activity along the visual pathway, resulting in an overall similar experience within some bounds. Of course, just like the discrete approach, it does not offer any insight as to why color is not interpolated. It is tempting to speculate, that this derives in some way from the mechanisms at the root of color constancy, which bias against objects changing color. This bias (prior) might be overruled, as motion processing networks reach an interpretation of the input sequence as a moving object. Ultimately, only imaging at the sub-second time scale of several cortical areas at the columnar resolution might shed light on the dynamics leading to the perceived abrupt color switch. In any event, given the lack of empirical data this phenomenon on the face of it does not seem to bear on the question of discreteness.

As the authors themselves note, other experimental results sometimes offered in support of discreteness can just as readily be explained by continuous accounts, citing examples such as the wagon-wheel illusion ([Levichkina et al. 2014](#)). Therefore, in lieu of conclusive experimental evidence for discrete perception that has yet to be provided, it is perhaps worthwhile to examine the plausibility of discrete perception on more theoretical grounds—it should be asked whether the neural substrates underpinning our perceptual capacities could in fact support discrete perception. We begin by examining the constraints placed at the network level, if discreteness is true, given the known properties of neuronal elements and systems.

Making Quick Work of the Cinematic Illusion

In support of discreteness, Herzog et al. note that a system cannot go beyond its (temporal) resolution. However, this observation cuts both ways: the rate at which the state of a neuronal network can switch is bounded by the temporal resolution constraints imposed by its elements—neurons, glia and receptors.

The reasoning is as follows: the manifestation of this time resolution can be thought as a (series of) low pass filter(s). Thus, if perceptual events switch too quickly, i.e. at a rate that falls entirely out of the neuronal bandwidth, they would manifest neurally as smooth changes. Further still, even if the rate is within the bandwidth, the faster the switches are, and therefore the closer to the upper limit of neuronal responsiveness, these switches will progressively become more graded. All this raises the question as to the significance of such “discreteness”.

To estimate what is in fact the case we analyzed voltage sensitive dye imaging data collected from an awake primate V1 during brief presentations of stimuli ([Fig. 2A](#)). These data were reported in [Omer et al. \(2013\)](#). They allowed us to estimate the impulse response (IR) function of the network at each cortical location within the imaged area (note that the population impulse response we find is consistent with single cell estimates—e.g. see [Yeh et al. \(2009, figure 1\)](#)). The IR can be used to linearize that system—i.e. use a linear model to predict the expected response of the network to arbitrary inputs. To test the capacity of V1 networks to support discrete state switches, we carried out a simulation in which the predicted population response of V1 to an ideal control signal switching at 2 and 20 Hz ([Fig. 2B, C](#)) was derived.

In [Fig. 2](#), it can be seen that even though primate V1 can support slow discrete state switches (e.g. at 2 Hz), discrete switching is not possible at rates that are plausibly within the range necessary for a discrete framework to explain the percepts associated with rapidly changing stimuli (e.g. at 20 Hz), and continuous responses are predicted.² Because, there is no reason to expect human V1 to be radically different in this regard, we suggest that V1 does not support discreteness in a phenomenologically meaningful way. As there is no reason to expect

2 Of course, if the system is highly nonlinear, the bounds suggested by our simulation might not be representative. Alas, the data analyzed here did not allow us to examine this question directly. However, a study by [Naaman and Grinvald \(2004\)](#), employing VSDI as well, addressed this question explicitly: oriented gratings were presented to anesthetized cats in immediate succession, and compared to the evoked responses to their presentation in isolation. It was found that response latency to the second stimuli was prolonged during successive presentation. To the extent these data are representative, they suggest that at least for some stimuli actual responses might be slower than the linear bound.

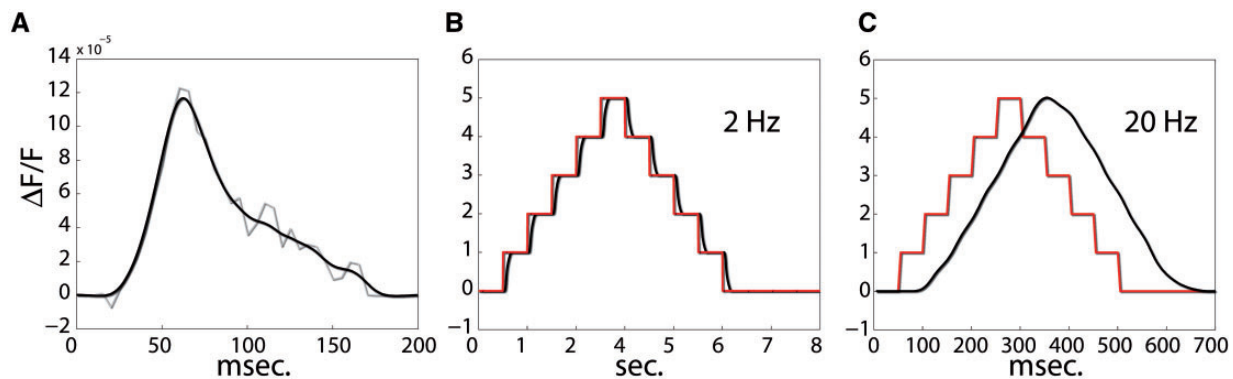


Figure 2. Visual cortex does not seem to support discrete switching at the perceptual sensitivity rate. (A) The impulse response function of V1 estimated through voltage sensitive dye imaging of an awake monkey. A primate was presented with brief (50 ms) presentations of oriented gratings. Responses were averaged, and FIR deconvolution was used to estimate the impulse response function (light grey; black—a smooth-spline approximation). (B and C) The impulse response computed in A can be used to derive a linear approximation of the expected response of the network to arbitrary inputs. Using ideal control signals (stair case modulations—marked in red) the network’s capacity for supporting discrete state switches at varying rates can be assessed. (B) The predicted response of V1 to state switches occurring at 2 Hz exhibits discrete switching (black). The control signal is superimposed in red. (C) The predicted response of V1 to state switches occurring at 20 Hz, a rate that is arguably within the rate at which perceptual content is modulated during perception of movement, shows no trace of state switching.

higher-order areas to be more sensitive temporally, this seems to strongly undermine the case for discreteness.

The proponents of discreteness could try and dodge the bullet by arguing that this limit could be overcome by having segregated neuronal populations implement state transitions, thereby circumventing such inherent limitations. Alas, not only is there overlap in the response properties of sensory neurons (non-zero width tuning), but there is mounting evidence that the population response to sensory events within sensory cortices are travelling waves (e.g. Muller et al. 2014; Alexander et al. 2013), further undermining this line of argument.

In the Interim

While presenting their framework, Herzog et al. focus on several experimental results. These examples, however, share a common characteristic—serial presentation of discrete stimuli presented in near isolation. Such focus runs the risk of glossing over an important distinction—namely between perceptual events, and what we will, for lack better terms, call total perceptual state. Even in the most tightly controlled experimental setting the total perceptual state is typically more expansive than the stimulus evoked perceptual events—comprising events and objects from several modalities at once. Further still, even within a given modality, several independent streams of events can impinge upon the senses at a given time, potentially giving rise to corresponding streams of perceptual events. Given the necessity of this distinction, the term percept is ambiguous as it can refer to perceptual events as well as total perceptual states. Given that ultimately a theory purporting to explain the fundamental properties of perception must hold for all types of perceptual events and settings, it is important to consider the plausibility of the discrete framework in light of this distinction.

In conscious states such as alert wakefulness, conscious experience is constantly changing, even in near complete lack of stimulation (Solomon et al. 1961). Moreover, at any point in time there are several objects and/or events populating our experience (typically against a dynamic background). How does this state of affairs constrain the underlying brain activity? If the initiation of simultaneously experienced perceptual events were synchronous, brain activity would grow more synchronous in conscious

states. However, the exact opposite is found—desynchronized EEG is the hallmark of conscious states (Steriade and Llinás 1988), suggesting asynchronous update of perceptual content during rich experience. Thus, even if the entry of a new item/type of content to the “stage” is clearly demarcated in time (see Fig. 3A), such events would show temporal overlap in neural activity. If so, it could be asked, even if we were to assume that each resulting modulation of the concurrent experience is somehow discrete, could that be said for the multiplexing of many such events? Plausibly, given a sufficient density of such asynchronous overlapping events, could the associated stream of experience be said to be meaningfully different from continuous unfolding of experience? Rather it should be assumed that our system is operating under the same computational and representational framework across the board, thus one that can accommodate dense multiplexing of asynchronous updates of sensory content. Let us ask then, if the proposed mechanism for discrete perceptual updates is in fact plausible outside of the context of a serial stream of unitary well-spaced events.

Not That Attractive After All

How does the brain “know,” the authors ask, when unconscious processing is complete and needs be rendered into a percept. They suggest that one possibility is that settling into an attracting steady state (ASS) leads to a signal that renders information conscious. Another possibility is that the attracting steady states themselves are conscious ones.

What does it mean to attain an attracting stable state as a result of stimulus driven processing? Within a given brain network, such as a cortical area, this would mean maintaining a stable firing pattern over some discernable span of time. Could ASSs be contained within a single or a few target cortical areas? This is highly implausible given that (i) focal brain damage in cortex is not associated with general loss of awareness and (ii) if we consider for example visual percepts, they generically contain not only high-level features such as objects, configuration and large-scale motion, but also detail down to (presumably) the retinotopic level (whether it’s motion, color, boundary contours and so on). Higher order visual areas that could potentially encode the high-level descriptors posited by Herzog et al., do not

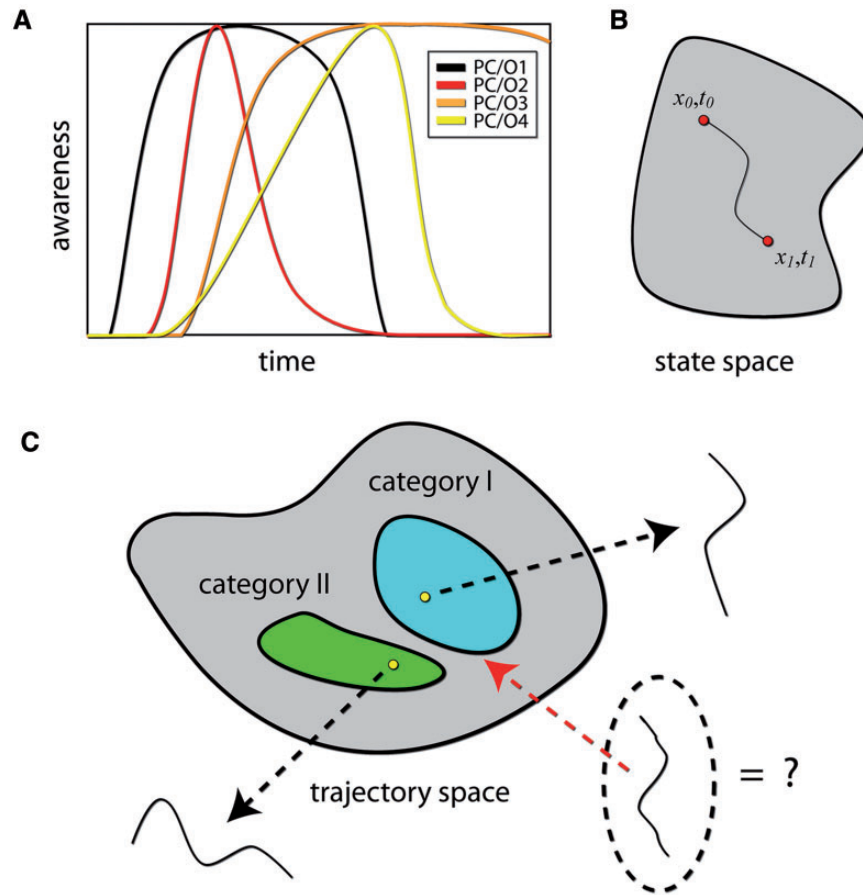


Figure 3. (A) The dynamics of perceptual content. Total phenomenological state comprises various articles whose temporal course of emergence and disappearance from awareness can vary greatly. Four perceptual contents/objects are shown (PC/O). Roughly speaking for a given stretch of time the total phenomenological state can be thought of as the sum of all the events/objects present in awareness. We have argued elsewhere that in a given state of consciousness this sum (the richness of consciousness, Tononi 2004) is constant over time (Fekete, 2010; Fekete and Edelman, 2011). (B) A trajectory through a state space. In the discrete case a trajectory would be a temporally ordered set of states, whereas in the continuous case it would be a path as is shown in the figure. (C) The space of all trajectories for a given system (in a given state of consciousness). Different categories correspond to different regions in space. Trajectories (points) within a category correspond to different instances of the category and share a family resemblance, through similarity in structure. This allows great flexibility in the face of varying uncertainty and computational demands, as the representational goal is achieved when a certain progression is achieved (a path), without implying that intermediate points need be reached at a constant rate, or that experience is strongly constrained by exact timing of stimuli. The ongoing hierarchical (multi-scale) computational efforts of interconnected neural networks enable the manifestation of conscious states, and constitute ever unfolding perceptual (and conceptual) content whose physical analog (i.e. isomorph) is emergence of ongoing multiscale structure (and synchrony) in space and time.

display differential signaling that is sensitive to nuance at the highest resolution of visual percepts. Therefore, the richness of percepts would require massive recruitment of both higher and lower level areas.

To attain an ASS spanning several cortical areas is mechanically challenging on several fronts: (i) There are considerable delays between cortical areas that are directly wired (e.g. up to 20ms in the ventral stream, Schmolesky et al. 1998). This would in effect shorten the window of opportunity such a massive coalition of neurons has to attain a discrete state switch in (in the stimulus driven case), as by hypothesis all areas instantiating a percept should attain the stable state within the same time window.³ Given that it is unlikely that single brain areas can

support a plausible state switching rate implied by experience the possibility will be categorically ruled out if requisite rate is, say, effectively doubled for higher level areas under rapid stimulation, (ii) The presence of different populations of neurons wired with differing latencies tends to be destabilizing (leading to oscillations, Jansen and Rit 1995) that manifest spatially as, e.g. traveling or standing waves, Fekete et al. 2017) and (iii) when the sensory apparatus is impinged upon by continuous/rapid streams of input, activity is destabilized, at least in lower sensory areas. If we consider movie viewing, where presumably activity can be modulated at rates such as 20Hz, the question must be asked how could ASSs be attained at a rate slower than

assume that the ASS destabilizes in the color network first during perception of the object—why assume that perception of the object as a whole is terminated then, and not that it is perceived as colorless (imagine a person who has a lesion in their color network)?

3 In fact, the assumption that not all areas taking part in an ASS rise and fall at the same time is in itself problematic. For example, if a percept requires the activation of a shape network and a color network,

the sensitivity (at the network level) of at least some lower sensory areas (especially during viewing smooth movement in which presumably there's a high degree of overlap in the underlying signaling networks).

The case for widespread ASSs is weakened even further by physiological data: if ASSs were the direct substrate of perceptual content update, we would expect to see pronounced synchrony during conscious states in comparison to say anesthesia, in which perception is abolished. However, as noted before, the exact opposite is observed—electrophysiology associated with a high degree of consciousness is marked by reduced global synchrony.

Alternatively, if an attracting steady state leads to a consciousness biomarker—which the authors admit that is currently unknown—it cannot be an ASS in itself as we have just seen that this does not solve the problem. Therefore, the biomarker must be some dynamical signal propagation. If that is the case, it is unclear why this would imply discreteness for perception. Rather, contents may enter our phenomenological stream at definite times, and exit the stage abruptly, an observation that would require fast initiation and termination of the neural processes involved, which is not at odds with continuous shifting of content while it is “on the stage” (Fig. 3A).

It could be asked if more general types of attractors, rather than ASSs, might be better suited for supporting discrete switching. To enable us to conclude that our analysis applies regardless of the particular dynamics following the switch, we need to consider in general when a measurable signal can be said to be discrete. In the ideal scenario this necessitates that the transition between two distinct successive states will be instantaneous. The reasoning is straightforward—the more sluggish the switch, the less separable the two states are, making the distinction between the states (and the time of their transition) an arbitrary decision imposed by an external observer rather than a matter of fact. Thus, regardless of the nature of the proposed states, for the neuronal case for discreteness to be plausible, at least some neurons should switch state nearly instantaneously (i.e. at a rate of change that is qualitatively different from non-switching epochs). This, according to common wisdom, would imply a near instantaneous change in the firing of neurons in a widespread network. In other words, our arguments against ASSs would apply to the letter in these scenarios as well.

Conclusion—Maybe Some Other Time

To date there is no direct empirical evidence favoring a discrete view of perceptual experience—all existing results are at the least consistent with a continuous outlook. On the contrary, one must ask what would be the significance (or meaning) of experience being discrete in some sense given that: (i) the temporal sensitivity of perception is not well removed from the limits of the temporal sensitivity of neural elements, therefore the system cannot meet the constraints implied by the latter while upholding fast discrete switching of perceptual content, (ii) the unrelenting multiplexing nature of experience in the awake state(s), which leaves no room “on the stage” for truly discrete events and (iii) our current understanding of neural computation and physiology militates against the implementation of (global) discrete events in any meaningful way.

Let us return therefore to the initial question posed by Herzog et al.: how could the apparent smoothness of experience be reconciled with erratic brief sensory sampling, and long integration times, together with much higher rates of change of perceptual content? Perhaps, it is instructive to start with the nature of experience itself: conscious states give rise to an

unfolding ever-changing coherent reality, populated by objects moving relative to dynamic backgrounds. If this end product is seen as the goal our system is trying to achieve, namely generating a coherent virtual reality, the representational processes enabling this should be up to the task.

Accordingly, a percept should be identified not with an end state, but with the entire process associated with it (Spivey and Dale 2006): we don't understand a sentence when we finish reading it, but whilst reading it our understanding unfolds and evolves (and can surely linger and evolve afterwards). In technical terms, this process is a trajectory through the state space of the brain.⁴ Under this perspective, the fact that different aspects of such a process arise with different latencies is not problematic. Rather, just as brain activity is organized across various scales in space and time, instead of assuming any such level of organization is somehow privileged, it stands to reason that all such structure contributes in varying amounts to the multilayered organization of experience in (cognitive) space and time.

This neurodynamic outlook to perception (van Leeuwen, 2015) sits well with (a kind of) hierarchical predictive processing (Lee and Mumford 2003; Friston 2005): We contend that one of the fundamental computations carried out by brain networks is prediction, and that these efforts take place on various levels of organization of the brain, from cells to networks. At the cell, assembly, and network level, predictions are continuously made of the expected activity of cells, assemblies and networks. At the level of the entire brain these predictions take a more familiar form of objects, events and scenes (and concepts), which are inferred as the (hidden) causes of our percepts. These predictions are generated and evaluated in an approximate⁵ hierarchical Bayesian framework (i.e. these computations implement a generative inferential model).

Cortical networks (areas) are densely and almost always reciprocally coupled (Felleman and Van Essen 1991). At any given point in time, we suggest, each such area actively predicts the future activity of the areas it is wired to (with the temporal and spatial scales of predictions varying according to the functional “specialty” of the pertinent network). These predictions are propagated to connected areas, as well as error signals pertaining to concurrently received predictions.⁶ These error signals are reciprocal—not exclusively ascending (Heeger 2017) or descending (Mumford 1992) as suggested by others—thus

4 This naturally raises the question as to what makes some brain trajectories, and not others (e.g. under anesthesia), conscious. The reader is referred to Fekete (2010), Fekete and Edelman (2011), Fekete and Edelman (2012) and Fekete et al. (2016), where this question is discussed at length.

5 By approximate we mean that networks instantiate Bayesian-like strategies incorporating priors (prior knowledge) to weight conditional probabilistic estimates. The maximal weighted estimate, instantiates a multifaceted joint distribution—the interpretation (meaning) of the current situation in which the agent is placed. We do not however claim that such computations are strictly speaking optimal—as other fundamental processes also constrain network activity, e.g. homeostasis at the sub-cellular and cellular levels.

6 It has been suggested that the architectural differences between ascending and descending pathways stem from a fundamental difference in the computations feeding into each type of pathway (e.g. error signals are propagated via ascending pathways, and predictions via descending pathways (Mumford 1992). We contend in contrast, that the difference stems from the convergent vs. divergent nature of ascending and descending pathways, e.g. as far as spatial and temporal receptive fields are concerned, which necessitates disparate connectivity patterns to maintain (suitably) balanced mutual influence.

potentially allowing the computational benefit displayed by predictive coding models, as well as the familiar “analysis by synthesis” (Heeger 2017; Clark 2013). This exchange is carried out without cessation at various temporal and spatial scales. As both the body and environment are typically moving/changing, the priors (weighting possible explanations by prior knowledge) on such prediction lean heavily towards change. This results in an elaborate dance in which mutual constraints are placed on the participating elements and networks.

From this perspective neural nets do not need to converge to a predication for a percept to emerge. Rather, in conscious states, mass orchestrated activity is a constantly unfolding prediction (explanation) that needs no further synchronization. The concurrent laborious efforts of the underlying network dynamics constrain the unfolding of the upcoming perceptual (and more generally phenomenological) changes in definite ways.

Indeed, collective behavior can exhibit scale free correlations. This is to say, dynamics enabling collective continuous coherent response to the environment, spanning distances that are independent of the spatial limits on interactions between individuals (Cavagna *et al.* 2010). This has been formalized using a brain-like connectivity model (the so-called connectome) as a dynamical system operating near criticality (Haimovici *et al.* 2013). Analysis of resting fMRI data provides direct evidence that brain networks are organized in such a scale free manner (Tagliazucchi *et al.* 2012), which can be thought of as dynamical compensation for the delay structure of the multitude of neural pathways imposed by neural signal conductance time.

This intrinsic dynamics of interdependent inferred hidden causes provides the background against which new sensory input is evaluated: Every level of the perceptual hierarchy provides a prediction, with which inputs are continuously matched to form prediction

errors, and although some of those prediction errors may be dampened (based on their relevance or expected uncertainty), perceptual continuity is preserved. Awareness cannot be attributed to any one level or region in the perceptual streams, and hence survives knock-out or interference in any one such region.

External stimuli should therefore be thought of as a perturbation that may cause the brain trajectory to change its course, provided that they pack enough punch (as different volumes in the space of possible trajectories are associated with different experiences (Shepard 1987; Edelman 1999; Gärdenfors 2004; Fekete 2010; Fekete and Edelman 2011), the greater the exerted pull, the greater the change in experience—Fig. 3C). Only evidence that violates our predictions in a significant way will actually cause noticeable shift in our percepts (relative to the prediction which is typically dynamic). What is significant will be very much dependent on attentional set, which is to say, the internal states holding active predictions and expected uncertainty or relevance of inputs. It will mean that substantial temporal or spatial discontinuities in inputs will often not sway our continuous experience (Koenderink *et al.* 2012). When in a conscious state, experience will unfold regardless of specific stimulation, because stimulation is no strict requirement for conscious experience given the predictive inferential machinery. Indeed, consciousness is not lost or diminished by say phasing external input out almost completely, as is the case when one is lost in thought, ruminating or daydreaming among other things.

The meshing of temporal (and spatial) scales as well as different modalities through ongoing hierarchical inference enables continuity in our experience. So we cannot state that there is a separate “visual consciousness,” though the visual (or any other sensory) contents might be very fragmented

(seemingly discrete) relative to other current contents of consciousness. Even though the visual contents of experience might be highly discontinuous, their experience is still continuous, given that it is never limited to this or that article. Rather, it is precisely because of the continuous background computational labor, that perceptual content can be perceived as discontinuous in various scales in space and time.

In sum—not only are the empirical results readily explained in a continuous framework, but the brain does not support virtual cinematography, and efforts to put it into a discrete straightjacket seem to offer hindrance rather than benefit. In other words, if your percepts appear continuous to you, perhaps it indicates that in this instance common sense does in fact make sense.

Acknowledgments

We wish to thank Amiram Grinvald and David B. Omer for providing us with the data reported here.

Funding

This work was partly supported by a Methusalem grant (METH/14/02, awarded to Johan Wagemans) from the Flemish government.

Supplementary data

Supplementary data is available at *NCONSC Journal* online.

Conflict of interest statement. None declared.

References

- Alexander DM, Jurica P, Trengove C *et al.* Traveling waves and trial averaging: the nature of single-trial and averaged brain responses in large-scale cortical signals. *Neuroimage* 2013;**73**:95–112.
- Cavagna A, Cimarelli A, Giardina I *et al.* Scale-free correlations in starling flocks. *Proc Natl Acad Sci USA* 2010;**107**:11865–70.
- Changizi MA, Hsieh A, Nijhawan R *et al.* Perceiving the present and a systematization of illusions. *Cogn Sci* 2008;**32**:459–503.
- Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci* 2013;**36**:181–204.
- Dennett DC, Kinsbourne M. Time and the observer: the where and when of consciousness in the brain. *Behav Brain Sci* 1992;**15**:183–201.
- Eagleman DM, Sejnowski TJ. Motion integration and postdiction in visual awareness. *Science* 2000;**287**:2036–8.
- Eagleman DM, Sejnowski TJ. The line-motion illusion can be reversed by motion signals after the line disappears. *Perception* 2003;**32**:963–8.
- Edelman S. *Representation and Recognition in Vision*. Cambridge, MA: MIT Press, 1999.
- Fekete T. Representational systems. *Minds Mach* 2010;**20**:69–101.
- Fekete T, Edelman S. Towards a computational theory of experience. *Conscious Cogn* 2011;**20**:807–27.
- Fekete T, Edelman S. The (lack of) mental life of some machines. In: Edelman S, Fekete T, Zach N (eds.), *Being in Time: Dynamical Models of Phenomenal Experience*, Vol. 88. Amsterdam: John Benjamins Publishing Company, 2012, 95–120.
- Fekete T, O’Hashi K, Grinvald A *et al.* Critical dynamics, anesthesia and information integration: lessons from multi-scale criticality analysis of voltage imaging data. Under review. 2017.
- Fekete T, van Leeuwen C, Edelman S. System, subsystem, hive: boundary problems in computational theories of consciousness. *Front Psychol* 2016;**7**:1041.

- Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1991;1:1–47.
- Freeman WJ. A cinematographic hypothesis of cortical dynamics in perception. *Int J Psychophysiol* 2006;60:149–61.
- Friston K. A theory of cortical responses. *Philos Trans R Soc B Biol Sci* 2005;360:815–36.
- Gärdenfors P. *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press, 2004.
- Haimovici A, Tagliazucchi E, Balenzuela P et al. Brain organization into resting state networks emerges at criticality on a model of the human connectome. *Phys Rev Lett* 2013;110:178101.
- Heeger DJ. Theory of cortical function. *Proc Natl Acad Sci* 2017; doi:10.1073/pnas.1619788114.
- Herzog MH, Kammer T, Scharnowski F. Time slices: what is the duration of a percept? *PLoS Biol* 2016;14:e1002433.
- Jancke D, Chavane F, Naaman S et al. Imaging cortical correlates of illusion in early visual cortex. *Nature* 2004;428:423–6.
- Jansen BH, Rit VG. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biol Cybern* 1995;73:357–66.
- Koenderink J, Richards W, van Doorn AJ. Space-time disarray and visual awareness. *Iperception* 2012;3:159–65.
- Kolers PA, von Grünau M. Shape and color in apparent motion. *Vision Res* 1976;16:329–35.
- Lee TS, Mumford D. Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 2003;20:1434–48.
- Levichkina E, Fedorov G, van Leeuwen C. Spatial proximity rather than temporal frequency determines the wagon wheel illusion. *Perception* 2014;43:295–315.
- Muller L, Reynaud A, Chavane F et al. The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nat Commun* 2014;5:3675.
- Mumford D. On the computational architecture of the neocortex. *Biol Cybern* 1992;66:241–51.
- Naaman S., Grinvald A. Dynamics of orientation state transitions in the visual cortex of the cat. Program No. 822.5. Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2004.
- Nadasdy Z, Shimojo S. Perception of apparent motion relies on postdictive interpolation. *J Vis* 2010;10:801.
- Omer DB, Hildesheim R, Grinvald A. Temporally-structured acquisition of multidimensional optical imaging data facilitates visualization of elusive cortical representations in the behaving monkey. *Neuroimage* 2013;82:237–51.
- Rao RP, Eagleman DM, Sejnowski TJ. Optimal smoothing in visual motion perception. *Neural Comput* 2001;13:1243–53.
- Scharnowski F, Rüter J, Jolij J et al. Long-lasting modulation of feature integration by transcranial magnetic stimulation. *J Vis* 2009;9:1–10.
- Schmolesky MT, Wang Y, Hanes DP et al. Signal timing across the macaque visual system. *J Neurophysiol* 1998;79:3272–8.
- Shepard RN. Toward a universal law of generalization for psychological science. *Science* 1987;237:1317–23.
- Solomon P, Kubzansky PE, Leiderman PH et al. Sensory deprivation. Cambridge, MA: Harvard university press, 1961.
- Spivey MJ, Dale R. Continuous dynamics in real-time cognition. *Curr Dir Psychol Sci* 2006;15:207–11.
- Steriade M, Llinás RR. The functional states of the thalamus and the associated neuronal interplay. *Physiol Rev* 1988;68:649–742.
- Stetson C, Cui X, Montague PR et al. Motor-sensory recalibration leads to an illusory reversal of action and sensation. *Neuron* 2006;51:651–9.
- Tagliazucchi E, Balenzuela P, Fraiman D et al. Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. *Front Physiol* 2012;3:15.
- Tononi G. An information integration theory of consciousness. *BMC Neurosci* 2004;5:42.
- van Leeuwen C. *Hierarchical Stages or Emergence in Perceptual Integration? Oxford Handbook of Perceptual Organization*. Oxford University Press Oxford, UK, 2015, 969–88.
- Van Wassenhove V, Grant KW, Poeppel D. Temporal window of integration in auditory-visual speech perception. *Neuropsychologia* 2007;45:598–607.
- VanRullen R, Koch C. Is perception discrete or continuous? *Trends Cogn Sci* 2003;7:207–13.
- VanRullen R. Perceptual cycles. *Trends Cogn Sci* 2016;20:723–35.
- Yeh C-I, Xing D, Shapley RM. “Black” responses dominate macaque primary visual cortex V1. *J Neurosci* 2009;29:11753–60.