



Vision Sciences Society Young Investigator Award 2011

Multiplexing in the primate motion pathway

Alexander C. Huk¹

Center for Perceptual Systems, Neurobiology, Psychology, The University of Texas at Austin, Texas 78712, United States

ARTICLE INFO

Article history:

Received 1 March 2012

Received in revised form 10 April 2012

Available online 20 April 2012

Keywords:

Motion

Multiplexing

3D

Encoding

Decoding

Interocular velocity difference

ABSTRACT

This article begins by reviewing recent work on 3D motion processing in the primate visual system. Some of these results suggest that 3D motion signals may be processed in the same circuitry already known to compute 2D motion signals. Such “multiplexing” has implications for the study of visual cortical circuits and neural signals. A more explicit appreciation of multiplexing—and the computations required for demultiplexing—may enrich the study of the visual system by emphasizing the importance of a structured and balanced “encoding/decoding” framework. In addition to providing a fresh perspective on how successive stages of visual processing might be approached, multiplexing also raises caveats about the value of “neural correlates” for understanding neural computation.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

This review explores the implications of a simple proposition: that any given neural circuit can potentially carry a multitude of signals. Such “multiplexing” is a core issue in engineering, and this review of the specific domain of visual motion processing suggests it may be a fruitful and illuminating concept to keep in mind in reverse-engineering the brain. I discuss this issue in the context of recent work on visual motion processing in the primate brain, but the concepts are intentionally general.

Many of us study vision because it allows for rigorous control of the inputs to the nervous system. The appeal of having such stimulus control is that it should afford more precise inferences about specific functions of neural circuits. With the ability to manipulate certain aspects of a computer-generated visual pattern while keeping others constant, we aim to isolate neural computations with a clarity that might not be as natural in stages of neural processing that are more distant from well-controlled sensory inputs.

Such inferential precision is a statement about the process of doing systems neuroscience, and not about how the brain must work. Although careful visual stimulus design has the potential to reveal and isolate neural computations, this does not logically imply that the neural computations themselves are implemented in a straightforward manner to grant us an easy process of discovery. If neurons and circuits can carry and combine multiple signals (and if later stages can then extract these signals), investigations of neural computations will benefit from considering multiplexing and demultiplexing as central issues.

The notion of individual neurons or circuits carrying multiple types of information is not new. In fact, this point derives from Rushton’s “Principle of Univariance”, which lies at the core of modern theories of sensory function and neural coding (Rushton, 1972). The Principle of Univariance states that the output of a sensory neuron, by itself, does not unambiguously signal a particular value of a single stimulus feature. Although originally stated in specific terms of the ambiguous mapping between photoreceptor output and both the wavelength and intensity of the light input, it can easily be generalized to many aspects of visual function. For example, the response of many single neurons in primary visual cortex (V1) is a function of both the orientation and the contrast of the visual pattern falling within the receptive field (among many other features; e.g., Dean, 1981). Thus, even if one knows the “preferred orientation” of the neuron (or has knowledge of the cell’s entire orientation tuning curve), a particular response level from that neuron could be the result of many combinations of orientation and contrast: An optimal orientation presented at a low contrast could produce a response identical to suboptimal orientations presented at higher contrasts, and so on.

The Principle of Univariance thus frames the output of neurons as potentially ambiguous, which in turn motivates consideration of “read-out” schemes that disambiguate these responses. For example, to explain how an unadulterated estimate of orientation is recovered from the responses of V1 neurons that are sensitive to both orientation and contrast, it is common to posit that multiple neurons with different orientation preferences are considered. By comparing the responses of multiple neurons, each with a distinct orientation preference, a read-out mechanism could disambiguate the effect of contrast (which affects all neurons similarly) to correctly arrive at the pure orientation.

¹ Elsevier/Vision Sciences Society Young Investigator Award 2011.
E-mail address: huk@mail.utexas.edu

Although this is a well-known pedagogical exercise in teaching vision science, its apparent simplicity may be misleading. In this particular example, the *encoding* of orientation and contrast appear very straightforward. Over the last few decades of research, however, the field learned that precise explanations of the interplay of contrast response functions and orientation tuning required an appeal to divisive normalization (Carandini & Heeger, 2011). This simple nonlinearity provided an elegant and general account of what initially appeared to be a set of unintuitive and complex empirical observations regarding contrast and orientation. Along the same logical lines, the *decoding* of one stimulus feature may also not be fully explained in this “cartoon” example: it is likely that the extraction of one piece of information (e.g., orientation) from a pattern of neural activity that carries many signals may also depend on computations that are not so intuitive. The challenge is how to unpack these more nuanced encoding and decoding computations in the face of richer (i.e., multidimensional) neural sensitivities.

A constructive path forward may derive from an explicit focus on “multiplexing”. In computer engineering, multiplexing is the passing of multiple signals through a common architecture (Hennessey & Patterson, 2011). This is a critical part of efficient circuit design and use, allowing a small amount of hardware to carry a variety of signals. Were it not for many clever forms of multiplexing, every signal would require its own dedicated hardware line, akin to early fixed line phone circuits that in fact required such a direct wired connection between caller and receiver (in fact, it would require two lines for bidirectional talk). Instead, modern computer circuits use a variety of tricks to use the same hardware for multiple signals. This requires an algorithm for multiplexing on one end, and a corresponding recipe for de-multiplexing (“demuxing”) on the other end.

Current conceptions of neural signaling do not, of course, assume the labeled line scenario of the earliest telegraph and telephone networks. Although most models of visual processing assume some form of distributed population representation, precisely what information is available in these representations (and how it is read out) is not well understood. Given the tremendous number (and flexibility) of sensory, cognitive, and motor processes instantiated in the primate nervous system, it seems not just possible, but probable, that multiplexing plays out in the brain. Next, I explain a potential instance of multiplexing drawn from recent work on motion perception, done by our lab and others. Although many of the details here remain speculative, this exercise suggests it will be fruitful for the multiplexing framework to guide the investigation of neural signals and circuits.

2. Co-existence of 2D and 3D motion signals in MT

Much is known about the visual processing of motion in the frontoparallel (“2D”) plane. In addition to the aforementioned sensitivities to orientation and contrast, many neurons in V1 exhibit a simple form of direction-selectivity (Movshon et al., 1985). Such “component motion” neurons signal the component of motion perpendicular to the cell’s preferred orientation. Many of these direction-selective neurons project to extrastriate visual area MT, where the vast majority (>90%) of cells exhibit compelling direction tuning that has been shown to directly relate to several aspects of the perception of motion (Albright, 1984; Parker & Newsome, 1998; Zeki, 1974a). Furthermore, many MT neurons are also disparity-tuned, which is independent of their direction tuning (DeAngelis, Cumming, & Newsome, 1998; DeAngelis & Newsome, 1999, 2004; DeAngelis & Uka, 2003). One might interpret this “separability” to imply that a given MT neuron carries information about a particular frontoparallel direction at a certain, fixed depth. Such

sensitivities are compelling and their mutual organization appears at first glance quite elegant. However, from an ecological perspective, this is puzzling: has the visual system evolved a brain circuit that encodes only frontoparallel (e.g., up, down, left, right, etc.) motions within fixed depth planes, but not motions towards or away from the observer? The more general lack of compelling evidence for 3D direction detectors in the primate brain was disconcerting, given that accurate perception of motions through depth is likely of central importance for the guidance of behavior.

Several years ago, we set out to make sense of the relations between the well-studied 2D motion system and the processing of “3D” motion (i.e., motions that contain significant components through depth). A brief review of these findings suggest that 3D signals may be multiplexed within the same circuitry known to process 2D motion signals. Although much of the evidence is indirect, the convergence of psychophysical and neuroimaging results have motivated us to perform a direct neurophysiological search for multiplexed 3D and 2D signals in single neurons, work which is currently in progress in our laboratory.

Historically, the study of 3D motion (or “motion in depth”, as it was often called) grew out of work on the static mechanisms of stereopsis, namely, binocular disparity processing (Julesz, 1960; Norcia & Tyler, 1984). It is therefore not surprising that the dominant framework for such “stereomotion” research focused on how the visual system might exploit changes in binocular disparity over time in order to extract 3D direction. But when an object moves towards or away from an observer, it casts moving images upon the two retina that have different horizontal velocities (Regan & Beverley, 1973; Regan & Gray, 2009). The simplest case is when a point (well approximated by your thumb) moves directly towards your nose: assuming fixed eye position, your right eye will “see” leftward motion of the point, while the left eye will “see” rightward motion. Although such an “inter-ocular velocity difference” (IOVD) is geometrically equivalent to the information available from changing disparity (CD), that does not imply that the visual system might process them equivalently. In fact, it seems likely that a CD mechanism arising from the building blocks of static disparity detectors would likely be quite different than an IOVD mechanism built from monocular motion signals. An even more basic question is whether the visual system encodes the CD cue, the IOVD cue, or both (Harris, Nefs, & Grafton, 2008). Fig. 1 schematizes the geometry of 3D motion and these two binocular cues.

Over the past four decades, the CD cue has enjoyed a stable amount of interest and empirical support for a role in 3D motion processing. One can generate a purely cyclopean 3D motion stimulus that contains the CD cue (by virtue of having steadily-changing disparities) but does not contain the IOVD cue (by virtue of not having coherent monocular motions to perform an inter-ocular comparison upon). This is done by “painting” a plane of fixed depth with a splatter of random dots. The disparity of that plane is then incremented (or decremented) gradually over time. But—from frame to frame of the video display—the exact locations of those dots are randomly replotted (Norcia & Tyler, 1984). This generates a compelling percept of a plane of TV-snow moving towards (or away) from the observer.

The IOVD cue, on the other hand, has been subjected to spottier interest over time, and less direct pieces of empirical support (Regan & Gray, 2009). This is primarily a geometric issue: it was long assumed that any stimulus containing an inter-ocular velocity difference must contain corresponding changing disparities (an assumption that will be discussed later). Instead, experiments assessing IOVD contributions to 3D motion processing focused either on a subtractive logic (i.e., comparing “Full cue” conditions that contain both CDs and IOVDs, versus CD-isolating stimuli), or using monocular motions per se (e.g., testing for relations between monocular motion processing and 3D motion perception).

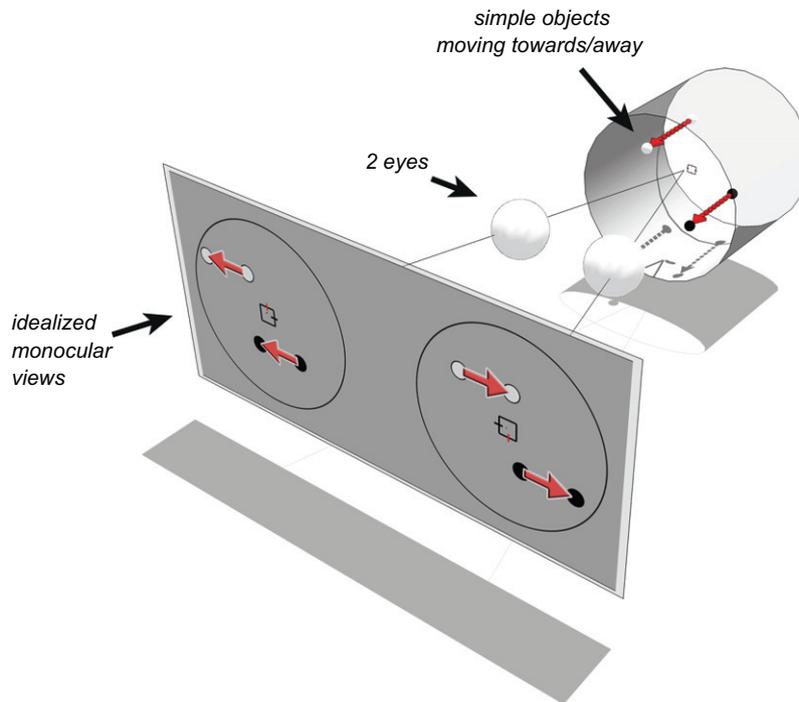


Fig. 1. Binocular viewing of 3D motion. When simple objects (like black and white spheres, top right) move towards and away from an observer (2 eyes indicated in center of scene), they project dynamic and distinct patterns upon the two retinæ (schematized on panels, lower left). Note that the disparity of the black and white dots changes over time (compare beginning and end of arrows), which corresponds to the “changing disparity” (CD) cue to 3D motion. Also note that the velocities of corresponding points have opposite horizontal directions (compare the left and right eye arrows), which corresponds to the “inter-ocular velocity difference” (IOVD) cue. When real objects move in the world, both cues are present in concert, but they can be dissociated experimentally (see text for details).

Our interest in the IOVD cue came about by accident. In generating some 3D motion displays using simple dark and bright dots on a grey background, we decided out of idle curiosity to see what a binocularly-anticorrelated display would look like. Such anti-correlated displays maintain the usual pattern of binocular pairings and disparities, except that dark elements in one eye are paired with bright elements in the other eye, and vice versa (imagine viewing a grayscale image in one eye, and its photographic negative in the other). Anti-correlation is well known to disrupt conventional disparity mechanisms in visual cortex, and in concert has devastating effects on psychophysical performance in depth-from-disparity tasks (Cogan, Lomakin, & Rossi, 1993; Cogan et al., 1995; Harris & Rushton, 2003). When we viewed the anticorrelated version of our 3D motion stimulus, we could immediately perceive that our ability to judge whether the dots were near or far was greatly impaired. This was no surprise. However, it was also very clear that we had no trouble at all discriminating whether the dots were moving towards or away from us. This was a compelling perceptual dissociation, and one that we took as evidence for the existence of a strong contribution of the IOVD cue: anticorrelation appeared (as usual) to reduce the strength or fidelity of the disparity signals that would be needed to compute the CD cue, so the IOVD cue was likely to be carrying the day. We performed a series of psychophysical measurements that confirmed this interpretation to be the case (Rokers, Cormack, & Huk, 2008), and went on to quantify the relative strengths of the IOVD and CD cues to 3D direction discrimination performance across a wide range of stimulus conditions (Czuba et al., 2010). In short, the IOVD cue was not only present, but appeared to be the primary perceptual cue.

These findings motivated us to think about how the IOVD might be computed in light of what we know about motion processing. We were forced to contemplate multiplexing head-on: either this 3D motion signal was computed in the well-studied circuits known to extract 2D frontoparallel motion, or there was a distinct “3D

motion circuit”. Although the notion of a distinct 3D motion circuit or area is not parsimonious, the possibility of multiplexed 2D and 3D processing seemed unlikely given that the majority of motion-responsive neurons in key visual motion areas (like MT) were already known to have compelling roles in frontoparallel motion processing—and tests for 3D motion selectivity had only confirmed the known sensitivities to 2D motion and static disparities (Maunsell & Van Essen, 1983; but see Zeki, 1974b). If 3D motion really was computed in the same circuits that were known to process 2D motion, it seemed likely that these two types of motion signal were multiplexed.

Seeking to resolve this issue, we performed additional experiments to more directly test whether the canonical motion pathway might also carry 3D motion signals. We reasoned that 3D motion signals might be best evoked by presenting stimuli that were optimized to support 3D motion percepts, as opposed to starting with adaptations of static disparity displays. Such 3D-motion-centric stimuli are also very different from the usual 2D frontoparallel motion stimuli: they are of course dichoptic, motions are horizontal (and opposite between the two eyes), and perhaps most critically, the retinal velocities for 3D motion are slow compared to those used in 2D motion studies. This is simply a geometric consequence: the projected retinal velocity is merely the frontoparallel component of the 3D motion trajectory. (For an intuitive extreme example, consider the case of a point moving directly towards the center of one eye: its retinal velocity will be zero).

A series of experiments provided convergent evidence that the middle temporal and medial superior temporal areas (MT and MST) are likely key stages of 3D motion processing, as they are in 2D motion. In a suite of psychophysical experiments, we found that the inter-ocular velocity difference can be computed using signals that are tempting to contemplate as “eye specific” versions of the sorts of motion signals seen in extrastriate areas like MT and MST—and not as the consequence of comparing monocular motion

signals of the sort seen in V1 (Rokers et al., 2011). It is classically established that neurons in V1 have generally small receptive fields and encode directional signals that are effectively one-dimensional, registering the component motion perpendicular to each cell's preferred orientation. Neurons in areas like MT are thought to integrate these ambiguous 1D signals over space—thus building larger receptive fields—and also over a range of component motions that are all consistent with a particular 2D pattern motion (Albright, 1984; Movshon et al., 1985). Such “pattern motion” neurons can be thought of as taking the 1D motion signals coming from component motion neurons to compute 2D direction (Rust et al., 2006). It should also be noted that canonical conceptions of the visual hierarchy contain monocular neurons in V1, but assume that such eye-specific channels have been merged into a single “cyclopean” stream in extrastriate cortices.

This canonical architecture sets up an intriguing pair of competing hypotheses for the computation of 3D motion based on IOVDs: 3D velocity is either extracted by comparing monocular 1D signals from V1, or some sort of hitherto-unidentified “eye-specific” 2D motion signals are present in extrastriate areas like MT. Both possibilities are perplexing. Either 3D motion is extracted directly from 1D components, leaving the classical 1D to 2D transformation separate; or 3D motion is extracted by comparing eye-specific 2D motion signals, despite the belief that such 2D signals are extracted after the point of binocular combination, at which point information about the eye-of-origin of signals has been discarded. To test between these possibilities, we designed a “dichoptic pseudoplaid” stimulus that could only support 3D motion percepts if the IOVD mechanism had access to eye-specific versions of 2D motion signals extracted over relatively large portions of the visual field. If the IOVD mechanism instead relied on the classical and well-established monocular 1D signals from V1, the stimulus would simply appear to be a jumble of monocular elements with no coherent 3D motion.

The stimulus comprises fields of small drifting gabor elements that are intentionally unpaired across the two eyes, but which are constructed to specify a single monocular pattern (2D) motion. This can be most intuitively understood in a cartoon example: imagine viewing a striped animal (such as a zebra or a white tiger) as it moves towards and away from you (i.e., a real object with a complex spatial pattern, engaged in 3D motion). Then, add the fact that you are viewing this spatiotemporal pattern from behind a structure that differentially blocks the two eyes' views (e.g., a dense bush). If this occluder granted the left and right eyes small glimpses of the pattern that were completely non-overlapping, your visual system would be challenged with the task of combining unpaired binocular information to extract 3D motion. A cartoon to support these intuitions is shown in Fig. 2A.

This is essentially what we implemented in a laboratory stimulus (Fig. 2B). To generate motion towards the observer, we defined a leftward retinal motion for the right eye, and a rightward retinal motion for the left eye. But each of these eye-specific pattern motions were actually instantiated by a sparse field of small drifting gabors with random orientations. In the right eye, all the speeds of the gabors were constrained to be consistent with a single leftward velocity; and vice versa for the left eye. (Again, this geometry is akin to viewing a translating object with a complex spatial pattern through a small set of apertures.) The critical bit of dichoptic geometry is that the locations of the gabors in one eye did not match the locations of the gabors in the other eye: in fact, we forced gabors in one eye to be at least 1.4 deg (edge-to-edge) from any gabors in the other eye. Because the classical receptive fields of V1 neurons are known to be smaller than that at the range of eccentricities of our stimulus, we presumed that any direction-selective V1 neurons had, at most, a single monocular gabor element within its receptive field (Van Essen, Newsome, & Maunsell, 1984).

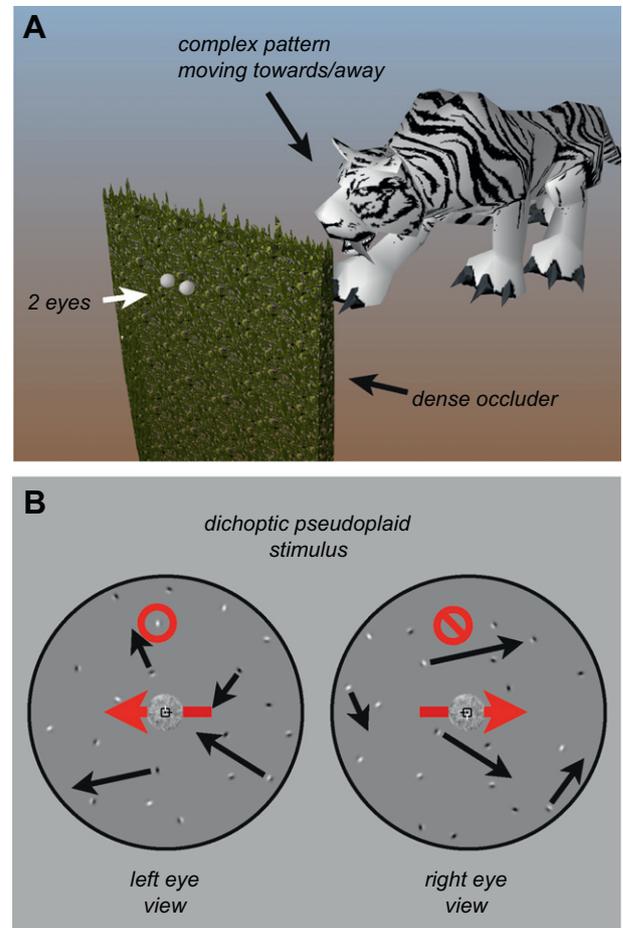


Fig. 2. Binocular viewing of partially-occluded 3D motion of complex patterns. (A) Cartoon for gaining intuitions for why IOVDs might be computed using eye-specific pattern motions. A scene involving an object with a complex pattern, viewed through a dense occluder, could yield binocularly-unpaired views of small patches of the object. Even in the absence of conventional binocular matching, it would be beneficial to compute IOVD-based 3D motion from the global pattern motions in each eye. (B) Schematic of the laboratory stimulus used to test for 3D motion percepts. The left eye and right eye contain opposite horizontal pattern motions, but these are supported only by a sparse set of gabors, each with a random local orientation. Critically, the gabors are spaced to be unpaired between the two eyes, without overlap on the scale of conventional V1 receptive fields. See text for more details.

Given this geometry, one might expect the cyclopean percept resulting from such dichoptically-separated stimuli to simply be a field of drifting gabors with varied orientations and velocities. Perhaps the percept would be further degraded by the unpaired nature of the elements, creating rivalry or a vague sense of binocular mismatch. To the contrary, when asked to perform a 3D direction discrimination task, subjects were able to perform well above chance, and exhibited accuracies that depended cleanly on the cosine of the 2D direction (i.e., on the horizontal component of the motion). This result is best explained by an IOVD mechanism that extracted a 2D direction for each eye by integrating the 1D gabor motion signals over many degrees of visual field. Such a sophisticated pattern-motion computation over large spatial regions seems very inconsistent with what we know about the function of V1, and instead motivates a search for such a mechanism in extrastriate cortex.

Indeed, in a related series of fMRI experiments, we found more direct evidence for 3D motion selectivity in extrastriate areas MT and MST in the human brain (Rokers, Cormack, & Huk, 2009). In one experiment, we found that MT and MST responded very strongly to dichoptically-opposite directions of horizontal motion

(which specify 3D motion) compared to dichoptically-opposite vertical motions, or to monocular-paired opposite motions of either orientation. In two other experiments, we found that MT and MST responded in distinct ways to stimuli that isolated the CD and IOVD cues, compared to corresponding control stimuli that contained the same building blocks (disparities and monocular velocities, respectively). Finally, we found that MT and MST exhibited direction-selective adaption to 3D motion that could be dissociated from the adaptation of early, monocular motion stages.

Although these are just the highlights, several other pieces of evidence from our group were also consistent with IOVDs relying on eye-specific versions of relatively late motion computations, and that 3D motion signals are dissociable from monocular and/or 2D constituents (Czuba et al., 2011; Czuba, Rokers, Huk, & Cormack, 2012). Together with work from other groups (e.g., Brooks, 2002, 2004; Fernandez & Farell, 2005, 2006; Nefs, O'Hare, & Harris, 2010; Sakano, Allison, & Howard, 2012; Shioiri, Saisho, & Yaguchi, 2000; Shioiri et al., 2008), these results paint a picture of a 3D motion circuit that uses eye-specific 2D motions as the key primitives. Of course, the obvious challenge is now to find direct neurophysiological evidence of 3D motion selectivity in extrastriate cortex—an effort we are already engaged in and one that we look forward to reporting soon. In the meantime, it is also enlightening to discuss why and how such 3D motion signals might have been hidden in prior attempts to find such selectivity.

The reason behind this may lie in the need to contemplate multiplexing when designing experiments. The most definitive test for 3D motion selectivity in single neurons of primate extrastriate cortex was performed by Maunsell and Van Essen (1983). They performed a rigorous set of measurements in monkey MT, measuring 2D direction tuning, assessing static disparity tuning, and then testing for 3D direction tuning. The take-away from this study was

that MT responses to 3D motion could be fully accounted for by separable contributions of 2D direction and static disparity tuning. The results are in fact so compelling that remarkably little work in primates has followed.

The conclusions drawn from this seminal study are certainly valid under the constraints of their stimulus set. Perhaps the most elegant aspect of the study was how the authors chose their set of 3D motion stimuli: they started by finding the preferred 2D direction of a given neuron, and then added subtle changes to the horizontal component of the velocity in one eye (leaving the optimal 2D direction in the other eye) to generate 3D directions. This is a clever way to slice the potentially-large stimulus space. It rests, however, on the assumption that 2D direction tuning would simply be the frontoparallel projection of 3D direction tuning. In essence, this would mean that the well-established 2D tuning of MT neurons was just a “flattened” version of underlying 3D tuning. And the results of this experiment compellingly demonstrate that this is not the case: probed in this way, MT responses simply reflected 2D tuning modulated by a preferred (fixed) disparity—no 3D tuning per se was needed. This raised the prospect that 3D motion was processed elsewhere (Likova & Tyler, 2007), effectively leaving MT as the “2D motion area”.

An alternate possibility is that MT neurons multiplex independent 2D and 3D motion signals. Instead of making the (reasonable) assumption that 3D and 2D tuning were geometrically linked in a direct manner, it is also possible that 3D tuning is wholly independent of 2D direction tuning. Our results to date suggest that the latter is a viable candidate for further consideration. Across our psychophysical and fMRI experiments, we generated stimuli from the perspective of simulating motion towards and away from the head, as opposed to taking more classical frontoparallel stimuli known to drive neurons effectively, and then adjusting these

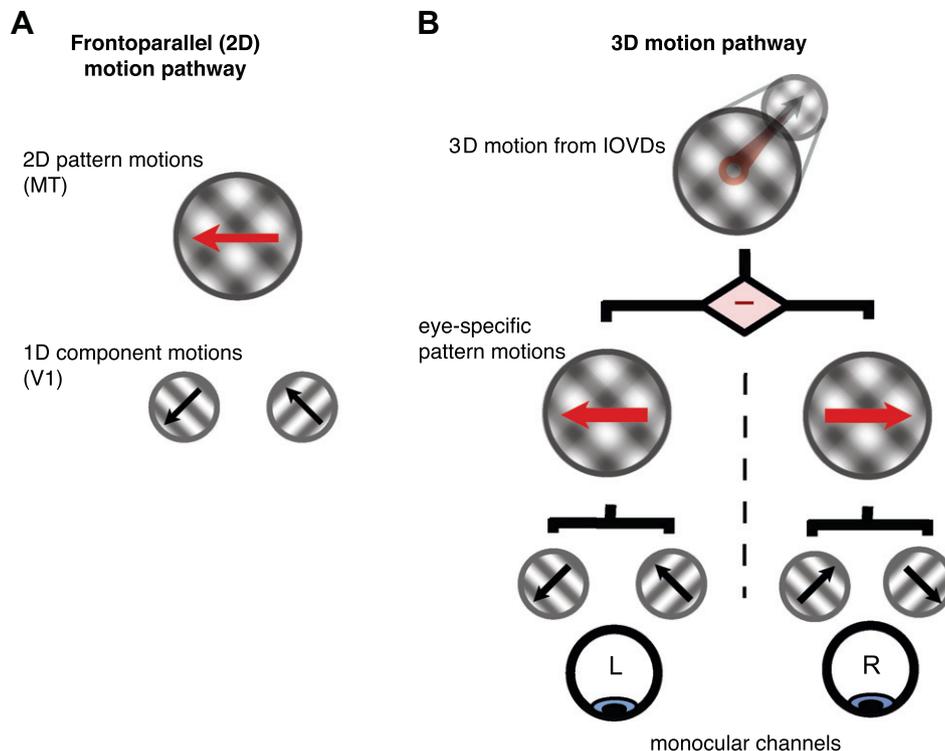


Fig. 3. Hierarchical models of 2D and 3D motion processing suggest multiplexing. (A) Conventional 2D (frontoparallel) motion system. Component motion neurons in V1 extract local 1D velocity estimates, and then pattern motion neurons in MT integrate these signals to compute the 2D velocity. (B) Putative 3D motion pathway. Early monocular channels (likely in V1) extract 1D component motions, which are then integrated into eye-specific 2D pattern motions. We have suggested that such eye-specific pattern motions involve area MT, despite standard assumptions that information regarding eye-of-origin has been discarded at the point of (earlier) binocular combination. Then, 3D motion (from inter-ocular velocity differences) is computed upon these eye-specific pattern motions. These computations may be implemented in the same circuits known to process 2D motion (panel A).

stimuli to contain motion through depth. The result is a set of very different stimulus conditions, containing very slow retinal motions (on order of 0.5–2 deg/s in each eye), which are primarily horizontal and opposite in the two eyes. Such stimuli, based on classical results in MT, would seem at best suboptimal, and at worst ill-suited for driving MT (why present nearly-stationary stimuli to neurons in “the motion area”?). Canonical MT neurons prefer brisk velocities; a rule-of-thumb is 10 deg/s as a standard “MT peak” for speed (Nover, Anderson, & DeAngelis, 2005). MT is also known to exhibit opponency, a weaker response to opposite motions within a local patch of the visual field (Heeger et al., 1999). But despite using stimuli that contain two obvious suboptimalities, we have repeatedly been able to generate strong perceptual performance in 3D direction discrimination, and also to find evidence for a role of MT/MST in the underlying neural processing.

Indeed, some finer details in the MT literature suggest deviations from these canonical properties: some MT neurons are in fact strongly responsive at very slow retinal speeds that are more consistent with the retinal projections of motions primarily towards and away from the observer (Krekelberg, van Wezel, & Albright, 2006; Nover, Anderson, & DeAngelis, 2005; Palanca & DeAngelis, 2003). Furthermore, conventional 2D motion opponency appears to operate on a V1 scale (Majaj, Carandini, & Movshon, 2007), and in monocular (or eye-specific) pathways (Tailby, Majaj, & Movshon, 2010). Finally, one of the earliest studies in monkey MT (Zeki, 1974b) qualitatively suggested the presence of neurons tuned to opposite directions of motion in the two eyes. All of these wrinkles suggest the potential for additional neural computations in MT that are potentially aligned with stereoscopic, 3D processing—although the exact implications for multiplexing are not yet known.

In summary, the primate motion pathway contains well-established mechanisms for processing frontoparallel (2D) motion. Over the last few decades of intense investigation of the primate motion pathway, sensitivity for 3D motion remained surprisingly elusive. Although it is possible that the visual system evolved a “2D” motion system for viewing television and performing psychophysics, and a distinct “3D” system for interacting with the dynamic real world, a convergence of evidence (and parsimony) motivates a search for 3D sensitivity within the existing 2D pathway. The work reviewed above suggests that 3D motion signals may be multiplexed in the same structures, and perhaps the same cells, as are known to carry 2D direction signals. Fig. 3 shows schematics of how 2D and 3D motion signals may be organized in the same large-scale circuits. Direct neurophysiological tests of this motion multiplexing proposition are now in progress. Regardless of the answers ultimately arrived at, this exercise to date motivates the consideration of multiplexing as a general property of neural circuits and signals. Below, I discuss some of the broader implications of the multiplexing perspective.

3. Discussion

The most important general implication of multiplexing is that the maximal response of a neuron becomes far less important in thinking about the meaning of its signals. It is entirely possible that the sorts of 3D motion stimuli which we favor do not drive MT neurons to fire as many action potentials as more conventional 2D motion stimuli do. But unless one assumes a conventional rate code in which each neuron is part of a single “labeled line” architecture (defined by the peak of their tuning curve for a single stimulus feature), the sheer magnitude of response is not a central issue for understanding neural coding. Instead, what matters is whether the neuron encodes useful information about a particular stimulus feature that could be decoded by a later stage. In the example of motion processing, this could involve a neuron with robust 2D

direction tuning that also exhibits clear 3D direction tuning—even if it exhibited an overall stronger response to peak 2D directions over 3D directions, or if there was no obvious geometric relation between the 3D tuning and the frontoparallel direction tuning.

This viewpoint frames decoding (or “read-out”) of neural signals as more fundamental than the encoding of signals. Although the importance of understanding both encoding and decoding of signals is not new, an appreciation of multiplexing makes it clear that the vast majority of interesting neural computations are likely to be better thought of as challenges of read-out, as opposed to encoded representations running along labeled lines. If multiplexing is commonplace, then most circuits either carry multiple signals simultaneously, or at least have the capability to carry different signals at different times or under different conditions. Thus, the challenge in understanding neural circuits and signals is characterizing how some subset of that information is extracted. Although the classical emphasis on characterizing the encoding done by sensory neurons remains a critical component, a focus on decoding need not be applied only to “cognitive” functions directly related to behavioral responses. Rather, the entire visual system, often conceived of as a hierarchy of encoding stages, might be more fruitfully approached as a cascade of decoding stages (Lennie, 1998). Different streams of processing may partially de-multiplex signals that are generally useful for certain tasks compared to others, but the signals at different levels of the hierarchy are likely still very high-dimensional.

There is not a clear line between multiplexing and the concept of coarse population codes. Although a coarse population code typically assumes that individual neurons likely contribute to representations of many different external stimuli (i.e., as suggested by the classic combinatorial arguments against “grandmother cells” in the ventral stream), multiplexing per se makes the more specific assertion that each of those neurons carries potentially distinct signals, and not just multi-purpose constituents of a population code. Furthermore, for the concept of multiplexing to be useful in guiding future hypotheses in the domain of neural computation, it is probably best to reserve the concept for signals that result from significant computation, as opposed to the sort of generic “building block” signals often observed at early stages of sensory transduction. At an extreme, photoreceptors in the retina may qualify as “multiplexers” of every visual signal in a rather liberal sense, but the later-stage instances of multiplexing contemplated in this review (e.g., of 2D and 3D direction) seem more likely to spark novel approaches to understanding neural signals. That said, the encode-decode perspective advocated here has already yielded significant insights into retinal processing (Pillow et al., 2008).

This review has drawn admittedly imprecise connections between “neural multiplexing” and proper multiplexing in computer engineering. But some of the more precise ideas from engineering could (and should) be adapted as starting points for read-out computations. The notion of frequency-division multiplexing has already entered the realm of putative computations in a variety of neural systems (Ballard & Jehee, 2011; Koepsell, 2010; Panzeri et al., 2010). It seems likely that other forms (e.g., time-division multiplexing, as well as spatial versions of multiplexing) will also seed interesting lines of work on neural computation (Bridgeman, 1982; Cariani, 2004; Fotowat, Harrison, & Gabbiani, 2011; Friedrich et al., 2009; Rucci, 2008; Seegraves, 2011). If one abandons the simplicity of single labeled-lines architectures and maximal firing rates, then the emphasis on demultiplexing read-out mechanisms beyond winner-take-all and vector-average—but which are still neurophysiologically plausible—becomes a significant area for new research that lags far beyond our understanding of the encoding side. Recently, analyses from machine learning like support-vector machines are becoming common in neuroscience, and can be thought of as a general-purpose statistical tool for demultiplexing

(Graf et al., 2011). Likewise, dimensionality reduction approaches from motor control may also offer more general insights into the mechanisms of decoding (Churchland et al., 2010). In a most relevant example, recent theoretical work has shown that a simple non-linearity can “unmix” eye-of-origin information from seemingly “cyclopean” signals well past the anatomical point of binocular combination (Lehky, 2011).

Multiplexing also brings perceptual learning to the forefront of important tools for understanding neural signaling. Far from being a niche topic, the improvements in task performance seen with practice can be thought of as reflecting a “tuning” of a demultiplexing scheme (Bejjanki et al., 2011; Huang, Lu, & Doshier, in press). This is anecdotally the case in the motion domain, as it is well established that learning to perform a speed discrimination task—and not being affected by irrelevant variations in temporal frequency or contrast—is a surprisingly challenging exercise for even experienced psychophysical observers (McKee, Silverman, & Nakayama, 1986). Furthermore, although subtle changes in sensory areas have been documented in many domains, brain areas more centrally implicated in read-out exhibit large changes during perceptual learning (e.g., Law & Gold, 2008). Although these are arduous experiments, it is exciting to see growing interest in this topic, and the resulting insights regarding how much of the visual system is better thought of as specialized encoding steps, versus more general encoding mechanisms followed by a plastic series of decoding stages.

Finally, if multiplexing is indeed rampant, then the status of explicit “neural correlates” should be reconsidered. Although much neurophysiological work seems focused on identifying single neuron responses that qualitatively mimic (and/or quantitatively account for) perceptual performance in a particular task, the multiplexing perspective raises two major caveats. First, the observation of a neural correlate under some condition does not logically imply that there is indeed a tight link between brain and behavior: given that neurons likely carry a multitude of signals, tests for neuron-perception (or neuron-behavior) correlations should aggressively manipulate other variables known to affect the neural response but which are irrelevant to the behavior (or vice versa). Second, there is no reason why an explicit neural correlate must exist at all: the encoding stages could represent a diverse bank of signals that must be judiciously read-out by motor planning circuits, leaving the links between neural signals and behavior better thought of as resulting from limits in the encoding and decoding over population responses that carry a wide array of information without forming an externally-obvious ‘representations’. Although both of these points are probably well-taken by many, the multiplexing perspective more strongly suggests that seeking neural correlates might not be a particularly fruitful primary goal, and that observations of such (cor-)relations might not be all that telling about neural computation.

Regardless of whether this argument against the importance of explicit neural correlates is relieving or controversial, the real strength of the multiplexing perspective is that it suggests a rigorous way forward. This “encoding/decoding framework” rests again on a loose appropriation of engineering concepts that are not in themselves novel—but taking this framework seriously can color future experimental design and data analysis. In multiplexing circuits, it will be critical to ask both encoding and decoding questions at each neural level. On the *encoding* side, instead of simply measuring tuning curves along a single stimulus dimension, it is necessary to measure responses to multiple combinations of stimulus features in order to identify and characterize multiplexing. This allows for one to build a full encoding model of what drives a neuron (at least given the range of stimuli entertained). Although this can reflect a combinatorial challenge, modern computing hardware should allow us to more judiciously choose our stimulus

sets and our sampling of parameters—the neurophysiological parallel of the psychophysical move from method-of-constant-stimuli to more efficient adaptive “staircase” techniques (Eyherabide, 2008; McManus, Li, & Gilbert, 2011; Watson & Pelli, 1983; Yamane et al., 2008). There are also subsequent analysis challenges associated with high-dimensional stimulus spaces and sensitivities, but again these are more practical than conceptual. Ultimately this could motivate a transition from classical analyses that derive from peri-stimulus time-histograms (PSTHs) conditioned on a particular stimulus feature, to multivariate characterizations of coefficients and/or kernels associated with multiple stimulus features. Such a framework can also be extended to estimate interactions between terms as well as canonical nonlinearities.

On the other side, it is also critical to assess how well external variables can be *decoded* from the output of the same neurons. If analyses of encoding focus on the mapping between a stimulus and the expected response of a neuron, the decoding perspective then grapples with how much information encoded by the neuron can be extracted from its noisy and multiplexed output. If a neuron only represented one stimulus feature, then encoding would be just the trivial mirror of the encoding, where the only added challenge is the noise associated with generating spikes on a particular trial. On the other hand, if a neuron multiplexes, the ability (of either data analysis, or the next neural step in processing) to decode a particular stimulus feature is nontrivial, and depends on how well the decoding step can demultiplex (Machens, 2010). Understanding the computations that support decoding could take a more central place in the study of mammalian visual cortical computations, and can be addressed in a variety of frameworks, from statistical optimality to neural and biophysical plausibility.

More generally, an explicit attempt to understand each step of neural processing from both encoding and decoding perspectives seems particularly ripe in the context of the primate motion pathway. This is because so much progress has been made on the encoding side, because the degree of multiplexing may be tractably limited, and because exciting instances of decoding links to behavior are already documented, such as that of “choice probability” in MT neurons (Britten et al., 1996). In short, there is lots more work to do, even in seemingly well-understood and apparently simple systems like the primate motion pathway.

Acknowledgments

I thank Elsevier and the Vision Sciences Society for the 2011 Young Investigator Award; and Jonathan Pillow, Lawrence Cormack, Alec Scharff, Leor Katz, Thad Czuba and Bas Rokers for comments on the manuscript. Thad Czuba also made critical contributions to the generation of the figures. It is a pleasure to thank the current and former members of my laboratory for actually doing the experiments that motivated the ideas in this article (especially Bas Rokers and Thad Czuba). This work and these ideas emerged in collaboration with Larry Cormack. This work was supported by NIH Grants R01-EY020592 and R01-EY017366, and NSF CAREER Award BCS-0748413.

References

- Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*, 52(6), 1106–1130.
- Ballard, D. H., & Jehee, J. F. M. (2011). Dual roles for spike signaling in cortical neural populations. *Frontiers in Computational Neuroscience*, 5, 1–12. <http://dx.doi.org/10.3389/fncom.2011.00022>.
- Bejjanki, V. R., Beck, J. M., Lu, Z.-L., & Pouget, A. (2011). Perceptual learning as improved probabilistic inference in early sensory areas. *Nature Neuroscience*, 14(5), 642–648. <http://dx.doi.org/10.1038/nn.2796>.
- Bridgeman, B. (1982). Multiplexing in single cells of the alert monkeys visual cortex during brightness discrimination. *Neuropsychologia*, 20(1), 33–42.

- Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebriani, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, 13(1), 87–100.
- Brooks, K. R. (2002). Interocular velocity difference contributes to stereomotion speed perception. *Journal of Vision*, 2(3), 218–231.
- Brooks, K. R. (2004). Stereomotion speed perception: Contributions from both changing disparity and interocular velocity difference over a range of relative disparities. *Journal of Vision*, 4(12), 1061–1079.
- Carandini, M., & Heeger, D. J. (2011). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62. <http://dx.doi.org/10.1038/nrn3136>.
- Cariani, P. A. (2004). Temporal codes and computations for sensory representation and scene analysis. *IEEE Transactions on Neural Networks*, 15(5), 1100–1111.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*, 68(3), 387–400. <http://dx.doi.org/10.1016/j.neuron.2010.09.015>.
- Cogan, A. I., Kontsevich, L. L., Lomakin, A. J., Halpern, D. L., & Blake, R. (1995). Binocular disparity processing with opposite-contrast stimuli. *Perception*, 24(1), 33–47.
- Cogan, A. I., Lomakin, A. J., & Rossi, A. F. (1993). Depth in anticorrelated stereograms: Effects of spatial density and interocular delay. *Vision Research*, 33(14), 1959–1975.
- Czuba, T. B., Rokers, B., Guillet, K., Huk, A. C., & Cormack, L. K. (2011). Three-dimensional motion aftereffects reveal distinct direction-selective mechanisms for binocular processing of motion through depth. *Journal of Vision*, 11(10), 18-1–18-18.
- Czuba, T. B., Rokers, B., Huk, A. C., & Cormack, L. K. (2010). Speed and eccentricity tuning reveal a central role for the velocity-based cue to 3D visual motion. *Journal of Neurophysiology*, 104(5), 2886–2899.
- Czuba, T. B., Rokers, B., Huk, A. C., & Cormack, L. K. (2012). To CD or not to CD: Is there a motion aftereffect from changing disparities? *Journal of Vision*, 12(4), 7-1–7-3.
- Dean, A. F. (1981). The relationship between response amplitude and contrast for cat striate cortical neurones. *The Journal of Physiology*, 318, 413–427.
- DeAngelis, G. C., Cumming, B. G., & Newsome, W. T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature*, 394(6694), 677–680.
- DeAngelis, G. C., & Newsome, W. T. (1999). Organization of disparity-selective neurons in macaque area MT. *Journal of Neuroscience*, 19(4), 1398–1415.
- DeAngelis, G. C., & Newsome, W. T. (2004). Perceptual “read-out” of conjoined direction and disparity maps in extrastriate area MT. *PLoS Biology*, 2(3), E77.
- DeAngelis, G., & Uka, T. (2003). Coding of horizontal disparity and velocity by MT neurons in the alert macaque. *Journal of Neurophysiology*, 89(2), 1094–1111.
- Eyherabide, H. G. (2008). Burst firing is a neural code in an insect auditory system. *Frontiers in Computational Neuroscience*, 2(4), 430–436. <http://dx.doi.org/10.3389/fncom.2008.10.003.2008>.
- Fernandez, J. M., & Farell, B. (2005). Seeing motion in depth using inter-ocular velocity differences. *Vision Research*, 45(21), 2786–2798.
- Fernandez, J. M., & Farell, B. (2006). Motion in depth from interocular velocity differences revealed by differential motion aftereffect. *Vision Research*, 46(8–9), 1307–1317.
- Fotowat, H., Harrison, R. R., & Gabbiani, F. (2011). Multiplexing of motor information in the discharge of a collision detecting neuron during escape behaviors. *Neuron*, 69(1), 147–158. <http://dx.doi.org/10.1016/j.neuron.2010.12.007>.
- Friedrich, R. W., Yaksi, E., Judkewitz, B., & Wiechert, M. T. (2009). Processing of odor representations by neuronal circuits in the olfactory bulb. *Annals of the New York Academy of Sciences*, 1170, 293–297. <http://dx.doi.org/10.1111/j.1749-6632.2009.04010.x>.
- Graf, A. B. A., Kohn, A., Jazayeri, M., & Movshon, J. A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature Neuroscience*, 14(2), 239–245.
- Harris, J. M., Nefs, H. T., & Grafton, C. E. (2008). Binocular vision and motion-in-depth. *Spatial Vision*, 21(6), 531–547.
- Harris, J. M., & Rushton, S. K. (2003). Poor visibility of motion in depth is due to early motion averaging. *Vision Research*, 43(4), 385–392.
- Heeger, D. J., Boynton, G. M., Demb, J. B., Seidemann, E., & Newsome, W. T. (1999). Motion opponency in visual cortex. *Journal of Neuroscience*, 19(16), 7162–7174.
- Hennessy, J., & Patterson, D. A. (2011). *Computer architecture: A quantitative approach* (5th ed.). Waltham, MA: Morgan Kaufmann.
- Huang, C.-B., Lu, Z.-L., & Doshier, B. A. (in press). Co-learning analysis of two perceptual learning tasks with identical input stimuli supports the reweighting hypothesis. *Vision Research*. <http://dx.doi.org/10.1016/j.visres.2011.11.003>.
- Julesz, B. (1960). Binocular depth perception of computer-generated patterns. *Bell System Technical Journal*, 39, 1125–1162.
- Koepsell, K. (2010). Exploring the function of neural oscillations in early sensory systems. *Frontiers in Neuroscience*, 4(1), 53–61.
- Krekelberg, B., van Wezel, R. J. A., & Albright, T. D. (2006). Interactions between speed and contrast tuning in the middle temporal area: Implications for the neural code for speed. *Journal of Neuroscience*, 26(35), 8988–8998.
- Law, C.-T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, 11(4), 505–513.
- Lehky, S. R. (2011). Unmixing binocular signals. *Frontiers in Human Neuroscience*, 5, 78. <http://dx.doi.org/10.3389/fnhum.2011.00078>.
- Lennie, P. (1998). Single units and visual cortical organization. *Perception*, 27(8), 889–935.
- Likova, L. T., & Tyler, C. W. (2007). Stereomotion processing in the human occipital cortex. *NeuroImage*, 38(2), 293–305. <http://dx.doi.org/10.1016/j.neuroimage.2007.06.039>.
- Machens, C. K. (2010). Demixing population activity in higher cortical areas. *Frontiers in Computational Neuroscience*, 4, 126.
- Majaj, N. J., Carandini, M., & Movshon, J. A. (2007). Motion integration by neurons in macaque MT is local, not global. *Journal of Neuroscience*, 27(2), 366–370.
- Maunsell, J. H., & Van Essen, D. C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology*, 49(5), 1148–1167.
- McKee, S. P., Silverman, G. H., & Nakayama, K. (1986). Precise velocity discrimination despite random variations in temporal frequency and contrast. *Vision Research*, 26(4), 609–619.
- McManus, J. N. J., Li, W., & Gilbert, C. D. (2011). Adaptive shape processing in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 108(24), 9739–9746. <http://dx.doi.org/10.1073/pnas.1105855108>.
- Movshon, J. A., Adelson, E. H., Gizzi, M. S., & Newsome, W. T. (1985). The analysis of visual moving patterns. In C. Chagas, R. Gattass, & C. Gross (Eds.), *Pattern recognition mechanisms* (pp. 117–151). New York: Springer.
- Nefs, H. T., O'Hare, L., & Harris, J. M. (2010). Two independent mechanisms for motion-in-depth perception: Evidence from individual differences. *Frontiers in Psychology*, 1, 155. <http://dx.doi.org/10.3389/fpsyg.2010.00155>.
- Norcia, A., & Tyler, C. (1984). Temporal frequency limits for stereoscopic apparent motion processes. *Vision Research*, 24, 395–401.
- Nover, H., Anderson, C. H., & DeAngelis, G. C. (2005). A logarithmic, scale-invariant representation of speed in macaque middle temporal area accounts for speed discrimination performance. *Journal of Neuroscience*, 25(43), 10049–10060.
- Palanca, B. J. A., & DeAngelis, G. C. (2003). Macaque middle temporal neurons signal depth in the absence of motion. *Journal of Neuroscience*, 23(20), 7647–7658.
- Panzeri, S., Brunel, N., Logothetis, N. K., & Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends in Neurosciences*, 33(3), 111–120. <http://dx.doi.org/10.1016/j.tins.2009.12.001>.
- Parker, A. J., & Newsome, W. T. (1998). Sense and the single neuron: Probing the physiology of perception. *Annual Review of Neuroscience*, 21, 227–277.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207), 995–999. <http://dx.doi.org/10.1038/nature07140>.
- Regan, D., & Beverley, K. I. (1973). Some dynamic features of depth perception. *Vision Research*, 13(12), 2369–2379.
- Regan, D., & Gray, R. (2009). Binocular processing of motion: Some unresolved questions. *Spatial Vision*, 22(1), 1–43.
- Rokers, B., Cormack, L. K., & Huk, A. C. (2008). Strong percepts of motion through depth without strong percepts of position in depth. *Journal of Vision*, 8(4), 6-1–6-10.
- Rokers, B., Cormack, L. K., & Huk, A. C. (2009). Disparity- and velocity-based signals for three-dimensional motion perception in human MT+. *Nature Neuroscience*, 12(8), 1050–1055.
- Rokers, B., Czuba, T. B., Cormack, L. K., & Huk, A. C. (2011). Motion processing with two eyes in three dimensions. *Journal of Vision*, 11(2), 10-1–10-19. <http://dx.doi.org/10.1167/11.2.10>.
- Rucci, M. (2008). Fixational eye movements, natural image statistics, and fine spatial vision. *Network*, 19(4), 253–285. <http://dx.doi.org/10.1080/09548980802520992>.
- Rushton, W. A. (1972). Pigments and signals in colour vision. *The Journal of Physiology*, 220(3), 1–31.
- Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, 9(11), 1421–1431.
- Sakano, Y., Allison, R. S., & Howard, I. P. (2012). Motion aftereffect in depth based on binocular information. *Journal of Vision*, 12(1), 11-1–11-15.
- Segraves, M. A. (2011). Signal multiplexing in neural circuits – The superior colliculus deserves a new look. *Frontiers in Integrative Neuroscience*, 5(5), 1–3.
- Shioiri, S., Nakajima, T., Kakehi, D., & Yaguchi, H. (2008). Differences in temporal frequency tuning between the two binocular mechanisms for seeing motion in depth. *Journal of the Optical Society of America A, Optics, image science, and vision*, 25(7), 1574–1585.
- Shioiri, S., Saisho, H., & Yaguchi, H. (2000). Motion in depth based on inter-ocular velocity differences. *Vision Research*, 40(19), 2565–2572.
- Tailby, C., Majaj, N. J., & Movshon, J. A. (2010). Binocular integration of pattern motion signals by MT neurons and by human observers. *Journal of Neuroscience*, 30(21), 7344–7349.
- Van Essen, D. C., Newsome, W. T., & Maunsell, J. H. (1984). The visual field representation in striate cortex of the macaque monkey: Asymmetries, anisotropies, and individual variability. *Vision Research*, 24(5), 429–448.
- Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33(2), 113–120. Springer.
- Yamane, Y., Carlson, E. T., Bowman, K. C., Wang, Z., & Connor, C. E. (2008). A neural code for three-dimensional object shape in macaque inferotemporal cortex. *Nature Neuroscience*, 11(11), 1352–1360. <http://dx.doi.org/10.1038/nn.2202>.
- Zeki, S. M. (1974a). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *The Journal of Physiology*, 236(3), 549–573.
- Zeki, S. M. (1974b). Cells responding to changing image size and disparity in the cortex of the rhesus monkey. *The Journal of Physiology*, 242(3), 827–841.