Visual Neuroscience (2013), **30**, 219–222.
Copyright © Cambridge University Press, 2013 0952-5238/13 $25.00
doi:10.1017/S0952523813000527

SPECIAL ISSUE

Linking Hypotheses in Visual Neuroscience: A Tribute to Davida Teller

PERSPECTIVE

Three comments on Teller's "bridge locus"

J. ANTHONY MOVSHON

Center for Neural Science, New York University, New York, New York

(Received October 31, 2013; Accepted November 1, 2013)

Abstract

The notion of a set of neurons that form a "bridge locus" serving as the immediate substrate of visual perception is examined in the light of evidence on the architecture of the visual pathway, of current thinking about perceptual representations, and of the basis of perceptual awareness. The bridge locus is likely to be part of a tangled web of representations, and this complexity raises the question of whether another scheme that relies less on geography might offer a better framework. The bridge locus bears a close relationship to the neural correlate of consciousness (NCC), and like the NCC may be a concept which is no longer precise enough to provide a useful basis for reasoning about the relationship between brain activity and perceptual experience.

Keywords: Perception, Linking hypotheses, Bridge locus

Introduction

In the early 1980s, Davida Teller, building on earlier work by Brindley (1960, 1970) and others, articulated a set of linking propositions that attempt to formalize the relationship between behavior and its underlying biological substrate in the context of vision and visual perception (Teller 1980, 1984; Teller & Pugh, 1983). Some of the propositions are so extensive in their scope that they are difficult to examine scientifically. But at the core are the empirically testable members of the family of linking propositions that posit specific links between brain activity and behavior, propositions "… to the effect that a specific mapping occurs between the states of particular, well-characterized visual neurons and particular perceptual states; that is, propositions of the form: human beings perceive X whenever the set of neurons Y is in the set of states Z" (Teller, 1984). Consideration of the full set of linking propositions is beyond the scope of this brief *Perspective*; here I concentrate on a particular feature of Teller's architecture of linking propositions, the notion of a "bridge locus," meaning in this context the identity and geography of the particular "neurons Y" that form the substrate for perceptual experience.

The idea of a bridge locus was clearly articulated by Teller and Pugh (1983):

> *Most visual scientists probably believe that there exists a set of neurons with visual system input, whose activities form the immediate substrate of visual perception. We single out this one particular neural stage, with a name: the* bridge locus. *The occurrence of a particular activity pattern in these bridge locus neurons is necessary for the occurrence of a particular perceptual state; neural activity elsewhere in the visual system is not necessary.*

Of course in pure form, this statement can be disputed only by dualists (Hart, 1996); the rest of us happily accept that neuronal activity forms the only basis for perceptual experience. But what kind of neuronal activity, in what kind of structure, and under what kind of constraints?

The idea of the bridge locus carries an implicit assumption that visual processing is serial and deterministic, as is made explicit by Teller and Pugh:

> *For the sake of this discussion we shall assume the visual system to consist in a sequence of deterministic maps indexed by k, k = 0, 1, 2, ... n, which maps the set of physiological states of the neurons at one stage to the set of physiological states of the neurons at the next stage, as follows:*

$$S \xrightarrow{M_0} \Phi^{(0)} \xrightarrow{M_1} \Phi^{(1)} \xrightarrow{M_2} \Phi^{(2)} \ \ldots \xrightarrow{M_n} \begin{array}{c} \Psi \\ \uparrow M^* \\ \Phi^{(n)} \end{array} \longrightarrow \ldots$$

> *M_0 maps the set S of external stimuli to $\Phi^{(0)}$, the set of quantum catch states of all photoreceptors. M_1, M_2, ... M_n maps the set of physiological states at one stage to the set of physiological states at the next stage. The symbol Ψ represents the set of all conscious visual perceptions. We explicitly assume the existence of a map M* that maps the set of physiological states $\Phi^{(n)}$ onto the set of visual perceptual states.*

Address correspondence to: J. Anthony Movshon, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003. E-mail: movshon@nyu.edu

https://doi.org/10.1017/S0952523813000527 Published online by Cambridge University Press

One may doubt whether the concept of a bridge locus is helpful if it is instantiated by a large and diffuse set of neurons, scattered across many areas of the brain. To provide a useful framework, the bridge locus containing *M** must be *compact* with respect to the rest of the visual system and the brain and flanked by neurons of contrasting function that both precede and follow it in a processing chain. The word "locus" carries the implication of distinct place and seems incompatible with the idea of a diffuse set of neurons scattered throughout the nervous system. The problem of localization—of finding the *M** map—is central to the discussions that follow.

The architecture of visual pathways

In the abstract, the notion of a bridge locus is clean and crystalline. There is one set of neurons that maps physiological states to perceptual states, and this set is both necessary and sufficient for perception. The first set of questions to arise about this idea is structural. Is the notion of such a locus plausible, given what we know about the anatomical structure of visual processing, especially the visual pathways of the primate cerebral cortex? Much of what Brindley and Teller write about linking propositions is essentially peripheral in conception—while the language is general, the thought is natural to the retina, the lateral geniculate nucleus, and the primary visual cortex, in the traditional view best represented by the serial feedforward scheme represented by Hubel and Wiesel (2004). When the linking propositions were first articulated, studies of visual processing outside the primary visual cortex were in their infancy, and the complexity of the architecture of the cortical visual system was only understood in its barest outlines. But this complexity is now known to be immense, and it has critical implications for the utility of the idea of the bridge locus.

The first comprehensive account of the connectional architecture of the extrastriate visual cortex (Felleman & Van Essen, 1991) emphasized the serial hierarchical nature of the cortical connectional matrix. The resulting map framed views of cortical visual processing for the first time as a set of parallel and serial interconnected streams of visual processing and encouraged the analysis of visual function in terms not of particular areas having particular functions, but instead in terms of functional streams of areas, richly interconnected and with overlapping functions. It soon emerged that this anatomically defined network of cortical areas has no unique hierarchical structure (Hilgetag et al., 1996), demonstrating that there is no evidence for a privileged or even preferred path of information flow through the network. For any given situation and stimulus, then, information might pass through the visual system over multiple parallel, serial, and recurrent pathways, a situation rendered even more complex by the likely contribution of cortico–subcortical loops which may complement or even dominate cortico–cortical signal flow (Guillery & Sherman, 2002). Even with these complications, the cortical network was thought to be relatively sparse, in the sense that many known areas were thought not to be directly connected—fewer than one third of the possible connections among the Felleman/Van Essen areas were known in 1991. But, as anatomical and computational techniques have improved, more and more connections among visual cortical areas have been discovered, and it is now clear that at least two thirds of the possible connections between areas exist, offering a staggeringly rich substrate for visual information to flow from multiple sources over multiple routes to multiple targets (Markov et al., 2012).

Where does this tangle of areas and connections leave the concept of a bridge locus? If there are many pathways through which information can pass from the visual system to the substrate of perception, what reason is there to believe in the existence of a single, stable unique map (*M** above) from physiological states to perceptual experience?

Probabilistic representations and Bayesian perception

The idea of a bridge locus implicitly carries with it a theory about how perceptual representations are organized. It is most compatible with the idea that the representation of a particular piece of sensory information is achieved by the activity of a small number of neurons for which that piece of sensory information is the "trigger feature" (Lettvin et al., 1959). The idea of a bridge locus is also in essence deterministic (Teller & Pugh, 1983). The idea is not incompatible with the noise present in all neural representations, but while noise can be included in the bridge locus framework, it is only as a nuisance and not as an essential feature of sensory coding. More recent approaches emphasize the role of stochastic population representations and make correspondingly less use of the concept of a deterministic network of neurons, each signaling its own trigger feature (Shadlen & Newsome, 1998; Averbeck et al., 2006).

Just as Bayesian approaches to statistical inference have become more prevalent in recent years, so Bayesian approaches to perception have become popular. Bayesian statistics are inferential and seek to bring all available evidence to bear on the question of what a statistician should believe about the true but unobservable state of the world that gave rise to a set of evidence. Because evidence always comes with a dose of uncertainty, the Bayesian approach seeks to estimate that uncertainty and use it to give different influences to different sources of evidence, most commonly a direct measurement and a set of prior experiences and beliefs. Bayesian approaches to perception offer a principled answer to a central problem: what an observer should believe about the world based on sensory evidence and historical knowledge.

The extension of Bayesian ideas to perceptual experience has its roots in Helmholtz (1924) and Gregory (1966), who argued that perception has much in common with statistical inference and scientific hypothesis testing, in that it constructs a representation of the world denoting a plausible hypothesis about reality, given the imperfect representation provided by our senses. In a Bayesian framework, this is akin to defining vision as perceptual inference, the estimation of the true structure of a scene from an image. The problem of perceptual inference in vision is ill-posed because the retinal image can be an arbitrarily complicated function of the visual scene and there is often insufficient information to determine the scene uniquely—the image on your retina might be of and look like your grandmother, but it could have arisen from an indefinitely large set of other arrangements of matter. The brain, or any artificial vision system, must make assumptions, and these assumptions are captured by Bayesian statistics in the form of distributions of prior probabilities, which capture the observer's knowledge or belief about the possible states of the world that might have given rise to a particular set of sensory data (Yuille & Bülthoff, 1996). These priors might arise through evolution or the accumulation of individual experience. The Bayesian approach to perception does not specify any particular representation for them, but it is compelling and clear that they exert a powerful influence on perceptual experience.

Some of the best evidence for Bayesian perception comes from visual illusions, which can offer an informative look at the way

https://doi.org/10.1017/S0952523813000527 Published online by Cambridge University Press

information from different sources mixes in brain processing. Gregory (1966) promoted the view that many illusions represent faulty hypotheses about the world based on imperfect sensory data. Applying a Bayesian approach to visual illusions has put this idea into a more robust formal framework. In this view, illusory percepts, although erroneous in the specific case that gives rise to the illusion, are nevertheless optimal in a general sense, given the uncertainty inherent in the visual input. Therefore, rather than thinking of visual illusions as a consequence of faulty processing, they can instead be interpreted as arising from a rational system designed to make optimal judgments in the face of uncertainty. That optimality is implemented by the development of a prior. For example, assuming that the visual system carries a prior (an implicit belief) that visual motion is usually both slow and smooth explains a variety of illusions (Weiss et al., 2002). Key questions in considering the notion of a bridge locus in a Bayesian world are where is the prior represented and how does it interact with sensory data (the likelihood function, if one is speaking in Bayesian) to compute the final percept (the posterior).

There are two popular ideas on how the prior is stored. In one view, a population of neurons encodes the prior in the same way that sensory-driven representations of likelihood are encoded. These representations then combine with the sensory likelihoods to yield perceptual outcomes, perhaps as simply as by addition of properly transformed signals (Jazayeri & Movshon, 2006). This kind of explanation can account for why the same scene can be perceived differently when observers are asked to make different judgments (Jazayeri & Movshon, 2007). Another view is that the prior is implicitly represented in the distribution of preferences and selectivities of neurons in the sensory representation (e.g., Girshick et al., 2011). Evidence that favors this latter interpretation includes showing that the distributions of the preferences and selectivities of neurons in sensory populations closely follow the pattern expected from a model of implicit representation (Ganguli & Simoncelli, 2011). It is plausible that both forms of representation exist, direct representations for priors that emerge and evolve dynamically from immediate experience, and distributional representations that emerge from longer time scales, including evolutionary ones, that influence the structure of perceptual representations.

So there is a good reason to accept that information from multiple measurements and prior beliefs can be represented by perceptual systems and blended in a statistically optimal way to produce perceptual outcomes. Signals from multiple sensory sources can be optimally combined (e.g., Ernst & Banks, 2002) and blended with evidence obtained from prior knowledge and belief (e.g., Weiss et al., 2002). In some cases, experiments have even uncovered neural circuits that represent and perhaps directly compute the statistically optimal representations that Bayesian theory predicts (Beck et al., 2008; Gu et al., 2008). All of this complexity does not by itself invalidate the concept of a bridge locus, but it does make explicit the requirement that it incorporate evidence from multiple sources and makes it likely that the neurons of the bridge locus are widely dispersed in space, time, and function.

The bridge locus and conscious awareness

The idea of a bridge locus is closely related to the question of the site of conscious awareness. Crick and Koch (2003) offer a framework for this: "… we are attempting to find the neural correlate(s) of consciousness (NCC) … In round terms, the NCC is the minimal set of neuronal events that gives rise to a specific aspect of a conscious percept." The NCC thus seems identical to Teller and Pugh's "*M**" map that translates physiological states to perceptual states. It is therefore not surprising that the quest for the NCC has been difficult for many of the same reasons that call the bridge locus into question, and that many areas widely dispersed in the brain have been offered as candidates for the NCC (Rees et al., 2002).

A case of interest is the role of extrastriate visual area MT (or V5) in visual motion perception—MT is a strong candidate for the bridge locus for motion and probably the strongest candidate for a cortical bridge locus in any domain of visual perception. The case for a particular role for MT is made by Block (2005), but Block's argument is more subtle than simply pinning the label "motion NCC" onto MT. For one thing, he argues for the existence of two kinds of NCC, one responsible for the *phenomenal* experience of motion and another responsible for the *access* of motion information to the brain's systems of categorization, reasoning, planning, and the control of action—MT is his candidate for the phenomenal NCC for motion. One challenge to this idea is experimental evidence that there are circumstances under which the experience of motion is directly counter to that predicted from the pattern of activity in MT (Hedges et al., 2011), suggesting that MT cannot be the only bridge locus for motion. But even in a more limited role, Block's proposal is not that MT always and inescapably translates neuronal firing into perceptual experience—he argues that reentrant circuits that carry recurrent signals back to MT must be active for MT to become the phenomenal NCC. This echoes the concern articulated earlier about the anatomical complexity of visual cortical pathways, but it also raises the disruptive idea that a bridge locus or NCC might serve that role only when certain patterns of activity are present and not otherwise. Others have debated that conscious awareness is associated not only with particular areas but also with particular temporal patterns of activity in those areas (von der Malsburg, 1981; Shadlen & Movshon, 1999). Introducing the requirement for the bridge locus of a particular temporal code or a particular activation pattern makes the idea both complex and conditional and therefore even more difficult to define and delineate.

Conclusions

Scholars of perception have long been addicted to "locus questions," and the challenge of addressing those is what first led Teller to formalize her ideas on linking hypotheses (Teller, 1980). In neuroscience, a huge premium is now placed on localizing particular functions in the brain. No conversation about perception, cognition, or action goes far without a consideration of where in the brain the function under discussion is "located." The tools of human brain imaging encourage us to think about the focal hotspots of activity associated with many kinds of processing. It is a short step from these to the identification of a bridge locus, but it is important to consider whether our predispositions and our tools are misleading us. Functional MRI is the favored tool of human brain imaging, but fMRI is well suited to detect focal highly active regions and far less effective at exposing patterns of weaker activity distributed across wide areas of the brain. This bias is enhanced by the practice of displaying data in maps showing statistical significance rather than signal strength and setting the statistical threshold high to show only the most active regions to avoid false positives (Lieberman & Cunningham, 2009). As a result, the most common tool of modern cognitive neuroscience exerts powerful implicit pressure to think about brain functions in terms of a small number of special places, each an area with a distinct and perhaps unique function, each,

https://doi.org/10.1017/S0952523813000527 Published online by Cambridge University Press

perhaps, a bridge locus. It is important to recognize this bias, and to realize that in a richly interconnected brain with countless paths for information flow, it is very likely that a large number of brain areas are critical for the representation of any particular function (Vezoli et al., 2004). In such a system, one must wonder whether the simple elegant concept of the bridge locus has been overtaken by the complexity of the nervous system. The problem in the end is not that the concept is false—it is not—but that the realities of brain architecture and representation make it too complex and amorphous to serve the purpose that Teller originally intended, to focus thinking on a crisp central dogma and to use that focus as a tool for discovery.

Acknowledgments

The thoughts in this essay evolved over many years and percolated through many conversations, none more important than the ones I was fortunate enough to have with Davida Teller. I regret that she did not live to argue them with me and to point out my conceptual errors, for which I must therefore accept sole responsibility.

References

Averbeck, B.B., Latham, P.E. & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews. Neuroscience* **7**, 358–366.

Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E. & Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152.

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Science* **9**, 46–52.

Brindley, G.S. (1960). *Physiology of the Retina and Visual Pathway*. London: Edward Arnold.

Brindley, G.S. (1970). *Physiology of the Retina and Visual Pathway* (2nd ed.). London: Edward Arnold.

Crick, F. & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience* **6**, 119–126.

Ernst, M.O. & Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433.

Felleman, D.J. & Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* **1**, 1–47.

Ganguli, D. & Simoncelli, E.P. (2011). Implicit encoding of prior probabilities in optimal neural populations. In *Advances in Neural Information Processing Systems (NIPS*10)*, ed. Lafferty, J., Williams, C., Zemel, R., Shawe-Taylor, J. & Culotta, A., pp. 658–666. Cambridge, MA: MIT Press.

Girshick, A.R., Landy, M.S. & Simoncelli, E.P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience* **14**, 926–932.

Gregory, R.L. (1966). *Eye and Brain: The Psychology of Seeing*. London: Weidenfeld & Nicolson.

Gu, Y., Angelaki, D.E. & Deangelis, G.C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nature Neuroscience* **11**, 1201–1210.

Guillery, R.W. & Sherman, S.M. (2002). Thalamic relay functions and their role in corticocortical communication: Generalizations from the visual system. *Neuron* **33**, 1–20.

Hart, W.D. (1996). Dualism. In *A Companion to the Philosophy of Mind*, ed. Guttenplan, S., pp. 265–267. Oxford: Blackwell.

Hedges, J.H., Gartshteyn, Y., Kohn, A., Rust, N.C., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. (2011). Dissociation of neuronal and psychophysical responses to local and global motion. *Current Biology: CB* **21**, 2023–2028.

Helmholtz, H.V. (1924). *Treatise on Physiological Optics* (trans. Southall, J.P.). Rochester, NY: Optical Society of America.

Hilgetag, C.-C., O'Neill, M.A. & Young, M.P. (1996). Indeterminate organization of the visual system. *Science* **271**, 776–777.

Hubel, D.H. & Wiesel, T.N. (2004). *Brain and Visual Perception*. New York: Oxford University Press.

Jazayeri, M. & Movshon, J.A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience* **9**, 690–696.

Jazayeri, M. & Movshon, J.A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* **446**, 912–915.

Lettvin, J.Y., Maturana, H.R., Mcculloch, W.S. & Pitts, W.H. (1959). What the frog's eye tells the frogs brain. Proceedings of the Institute of Radio Engineers **47**, 1940–1951.

Lieberman, M.D. & Cunningham, W.A. (2009). Type I and Type II error concerns in fMRI research: Re-balancing the scale. *Social Cognitive and Affective Neuroscience* **4**, 423–428.

Markov, N.T., Ercsey-Ravasz, M.M., Ribeiro Gomes, A.R., Lamy, C., Magrou, L., Vezoli, J., Misery, P., Falchier, A., Quilodran, R., Gariel, M.A., Sallet, J., Gamanut, R., Huissoud, C., Clavagnier, S., Giroud, P., Sappey-Marinier, D., Barone, P., Dehay, C., Toroczkai, Z., Knoblauch, K., Van Essen, D.C. & Kennedy, H. (2012). A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cerebral Cortex*, in press [Epub ahead of print]. doi:10.1093/cercor/bhs270.

Rees, G., Kreiman, G. & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews. Neuroscience* **3**, 261–270.

Shadlen, M.N. & Movshon, J.A. (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron* **24**, 67–77.

Shadlen, M.N. & Newsome, W.T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *The Journal of Neuroscience* **18**, 3870–3896.

Teller, D.Y. (1980). Locus questions in visual science. In *Visual Coding and Adaptability*, ed. Harris, C.S., pp. 151–176. New York: Lawrence Erlbaum Associates.

Teller, D.Y. (1984). Linking propositions in visual science. *Vision Research* **10**, 1233–1246.

Teller, D.Y. & Pugh, E.N., Jr. (1983). Linking propositions in color vision. In *Colour Vision: Physiology and Psychophysics*, ed. Mollon, J.D. & Sharpe, L.T., pp. 577–589. London: Academic Press.

Vezoli, J., Falchier, A., Jouve, B., Knoblauch, K., Young, M. & Kennedy, H. (2004). Quantitative analysis of connectivity in the visual cortex: Extracting function from structure. *Neuroscientist* **10**, 476–482.

von der Malsburg, C. (1981). The correlation theory of brain function. Max Planck Institute for Biophysical Chemistry, Internal Report 81–2. Reprinted in *Models of Neural Networks II*, ed. Domany, E., van Hemmen, J.L. & Schulten, K. Berlin, Germany: Springer.

Weiss, Y., Simoncelli, E.P. & Adelson, E.H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience* **5**, 598–604.

Yuille, A.L. & Bülthoff, H.H. (1996). Bayesian decision theory and psychophysics. In *Perception and Bayesian Inference*, ed. Knill, D.C. & Richards, W., pp. 123–161. Cambridge: Cambridge University Press.

https://doi.org/10.1017/S0952523813000527 Published online by Cambridge University Press