

Varieties of perceptual truth and their possible evolutionary roots*

Shimon Edelman

Department of Psychology

Cornell University

September 2, 2014

Abstract

Hoffman, Singh, and Prakash (2014) observe that perception evolves to serve as an interface between the perceiver and the world and proceed to reason that percepts need not, or even cannot, resemble their objects. I accept their premise, but argue that there are interesting ways in which perception can be truthful, with regard not to “objects” but to relations, and that evolutionary pressure is expected to favor rather than rule out such veridicality.

Introduction

In the four decades since the publication of (Dobzhansky, 1973), many of us in the cognitive sciences have invoked his title — “Nothing in biology makes sense except in the light of evolution” — in print and in classroom discussions. Typically, however, such invocations amount to little more than a flat “because” answer to a bothersome “why” question. In their remarkable paper, Hoffman, Singh, and Prakash (2014; hereafter HSP) advocate the use of evolution as the key guiding principle, not to mark the final resting place for a certain type of psychological inquiry, but rather to motivate a reassessment of the entire psychology of perception.

Adopting an analogy from jurisprudence, one may take the HSP paper to be a draft of the first article of a constitution, which would impose certain strictures on all manner of psychological explanations. To

*A commentary on Hoffman, D., M. Singh, and C. Prakash, *The interface theory of perception* (2014).

succeed, such a project must amount to more than insisting on a certain style of informal reasoning about specific phenomena (as in what came to be known as “evolutionary psychology”; Cosmides and Tooby, 1997) — it must offer a formal framework for developing and evaluating models and theories. This, HSP does.

While psychology at large has not yet been blessed with the kind of general quantitative laws that serve as the foundations of physics, the situation in perception science is more advanced. Explicit computational approaches to problems in perception have become the rule rather than an exception; ideal observer methods (Barlow, 1980; Knill and Kersten, 1991) are well-established; and there is even work on general laws — as in the “universal law of generalization,” derived by Shepard (1987) from first principles of probability theory and Bayesian reasoning (see Tenenbaum and Griffiths, 2001 for a more recent discussion). Interestingly, the theory that underlies Shepard’s law of generalization both embraces evolution and predicts veridicality — a combination that, according to HSP, should not be possible. In this commentary, I propose to take a closer look at veridicality and reexamine its supposed incompatibility with evolution.

What is truth?

The problem of veridicality is unique in perception in that on a fundamental level it is a central concern not just for perceptual psychology but for all science, insofar as it aims to discover truths about the world through observation and intervention. In this context, we may begin by considering the (probably inadvertent) move on the part of HSP to both eat their cake and leave it whole. If, as the target paper observes, “natural selection tunes perception to payoffs, not to truth” (pp.21-22) and if “perception is about having kids, not seeing truth” (p.31), what basis can there be for the authors to claim also that “the environment in which our species evolved is a highly structured place, containing many regularities” (p.33)?

One way to avoid the problem here is to treat the latter statement as a metaphysical stance, that is, to profess belief in the existence of certain regularities “out there” in the world — regularities without which science would be impossible — without attempting to ground those regularities in perception. Now, I do believe that a bit of explicit metaphysics early on can save a lot of trouble later (Edelman, 2011), but as a scientist I would rather have as little of it as possible. In the present case, it seems that the apparent

contradiction can be resolved if we clarify the concept of “truth” as it is used throughout the target paper.

What, then, is the meaning of “truth in perception”?¹ I find HSP’s statements in this connection entirely unobjectionable. In particular, I agree that there is no conceivable way for the taste of a substance to be “true” (p.61: “What can we possibly mean by the *veridical* taste of a molecule?”). Likewise, the neural representation of an object that is perceived as a car is merely some neural activity, and as such is nothing like what’s actually out there (p.69: “But the objective reality is not a car and doesn’t remotely resemble a car”). However, rather than ending the argument about the possibility of truth in perception, these observations should be seen as the starting point in a search for ways to make it work.

To move forward, we must determine whether perception can evolve to be truthful (thereby helping ground science more in physics, biology, and psychology, and less in metaphysics) in other ways than the one that the above quotes from HSP expose as absurd. At least three such ways suggest themselves. I shall first introduce them, then briefly discuss each possibility in turn, focusing on the question whether or not it can emerge and survive in an evolutionary setting.

- *Categorical consistency* — I call a perceptual channel (HSP, p.41), denoted by a function $f : \mathcal{X} \rightarrow \mathcal{Y}$, *CC*-truthful if, when given as input any member x of a class of stimuli \mathbf{X} , $x \in \mathbf{X} \subset \mathcal{X}$, it reliably² evokes a representation $f(x) = y \in \mathcal{Y}$. Here, \mathcal{X} and \mathcal{Y} are sets, each equipped with its own identity relation that defines set membership. Categorically consistent perceptual channels make it possible to perceive truthfully the identity (modulo irrelevant transformations and noise) between stimuli, as well as the recurrence or persistence of a stimulus over time. Example: recognizing the face of a familiar person as such, despite the usual variability in appearance.
- *Second-order isomorphism* — a *CC*-truthful perceptual channel is *SOI*-truthful if both \mathcal{X} and \mathcal{Y} are metric spaces and, further, if the application of $f : \mathcal{X} \rightarrow \mathcal{Y}$ results in a reliable and sufficiently consistent mapping of the rank order of similarities between members of \mathcal{X} to the rank order of similarities between their corresponding representations in \mathcal{Y} . Example: perceiving two shades of

¹A question that echoes the opening sentence of Francis Bacon’s essay *On Truth* (1601): “What is truth? said jesting Pilate, and would not stay for an answer.”

²This qualifier and others like it below can be made precise with a bit of extra notation.

yellow to be more similar to each other than to a shade of green (think ripe and unripe bananas).

- *Causality* — a *CC*-truthful perceptual channel is *Cau*-truthful if it reliably maps causal relations R_X over elements of \mathcal{X} to isomorphic causal relations R_Y over elements of \mathcal{Y} . (Causality is of course a notoriously difficult concept (De Pierris and Friedman, 2013; Schaffer, 2009; Lagnado, Waldmann, Hagmayer, and Sloman, 2007); for our present purposes, it suffices, in the spirit of Hume (1748, 4.6), to define R_X simply as a time-ordered pair of cause and effect $(x_1^{t_1}, x_2^{t_2})$, where $x_1, x_2 \in X$ and $t_2 = t_1 + \Delta t$ for some sufficiently small Δt , with R_Y defined accordingly.) Example: taking *thunder* to be caused by *lightning* (while possibly remaining perfectly ignorant as to what each of them “really” is).

Categorical consistency

Although this and the other two types of veridicality have been extensively discussed in the literature (see (Edelman, 2008) for an overview), much remains to be said. I shall therefore keep my remarks to an absolute minimum, focusing mainly on evolvability. With regard to *CC*-veridicality, we should note that the structure of categories depends on the choice of features and their use, as determined by fixing the perceptual channel f (for a formal argument, see, e.g., Watanabe, 1969, pp.376-377). This choice is precisely the “pressure point” through which evolution may exert its influence, by conferring selective advantage on perceivers who stumble on a good “prior spacing of qualities” (Quine, 1969; see also Clark, 1993).

A concept that can be used to link this type of evolutionary thinking to the psychology of categorization is that of *consequential region*, introduced by Shepard (1987) in the context of his proposed universal law of generalization. Given a stimulus that has already been encountered, with known consequences (say, an apple that has been found, pondered, tasted, and enjoyed), another comparable stimulus (a second apple) should be assigned to the same category as the first one (“crunchy,” as opposed to “mushy”) if it happens to fall within the consequential region defined by the first stimulus in the representation space. Shepard (1987) was able to reach some surprisingly deep and general results regarding categorization and generalization, by reasoning probabilistically from first principles about the shape of the putative consequential region associated with a stimulus.

Second-order isomorphism

The preservation of similarity ranks by a perceptual channel f is an instance of what Shepard and Chipman (1970) called the *second-order isomorphism* between (aspects of) the world and (the corresponding aspects of) its representation. On this account, first-order isomorphism obtains when a percept resembles its object — a nonsensical proposition, as pointed out by many theorists, including HSP. In contrast, second-order isomorphism requires that some relation over percepts resemble the corresponding relation over their objects. I have in the past referred to this distinction as the (first-order) representation *by* similarity, vs. the (second-order) representation *of* similarity (Edelman, 1998). The extensive available evidence for perceptual *SOI*-veridicality has been reviewed elsewhere (Edelman, 1999; Edelman and Shahbazi, 2012).

It is important to note that *SOI*-veridicality cannot be expected to hold for objects that differ too much from each other, that is, are too far apart in the representation space in the sense of Shepard (1987) — as, for instance, is the case for two hues that are diametrically opposite on Newton's color circle, such as orange and blue. Just how far apart the representations of two objects end up depends not only on their physical make-up but also on the perceiver's needs (witness, for instance, the categorical perception effect; Macmillan, Kaplan, and Creelman, 1977; Harnad, 1987), which are, of course, shaped by evolution (this makes the present account consistent with the spectrum example discussed by HSP). Locally, however, *SOI*-veridicality is not only possible, but is in fact generic: it obtains under any smooth perceptual channel function f (Edelman, 1999). I suspect that in terms of the analysis of the effect of group action offered by HSP, the global/local distinction corresponds to the distinction between classical and Lie groups (the latter mentioned on p.72 of the target paper as being outside the scope of the discussion).

As HSP note in a related context, the all-important questions here are, first, whether or not smooth perceptual channels become available often enough in the course of evolution, and, second, whether or not such channels survive competition with ones that are not smooth in the requisite manner. I conjecture that the answer to both questions is affirmative. The default initial condition for a newly acquired channel (corresponding to the emergence of a new sensory modality) seems to be inefficiency, such that it takes a large change in the stimulus to bring about even a small change in the response. This would make for a tightly knit representation space, where all things in the new modality feel alike. Evolutionary pressure would

then work to improve sensitivity (if it's worth the species' while; if not, the new modality will sink back into oblivion), so that eventually small changes in the stimulus bring about small changes in the response, making the channel smooth and therefore generically *SOI*-veridical.

Causality

The concept of *Cau*-veridicality introduces a new dimension to the representation space, that of time. Given the importance of serial order of events in guiding and forming behavior (Lashley, 1951; Kolodny and Edelman, 2014), temporal succession that is close enough (Lagnado and Speekenbrink, 2010) between salient $x_1^{t_1}$ and $x_2^{t_2}$ must be noted and the subsequent input monitored for possible predictive value of x_1 with regard to x_2 (see (Pearl, 2009) for a formal basis and (Holyoak and Cheng, 2011) for a review in the context of psychology).

Knowledge of the causal structure of the world, which is what this process aims to learn, is more useful than mere knowledge of the joint probability over the variables of interest (Pearl, 2001, 2013). Such knowledge can be made more reliable by taking special note of actions and by employing targeted interventions (which in science are called experiments) — a strategy that has been documented both in rats (Leising, Wong, Waldmann, and Blaisdell, 2008) and in human children (Kushnir and Gopnik, 2005). Causal knowledge, at least of simple two- or three-variable situations, is thus certainly easy enough to glean from temporal succession and intervention data. Could true causal knowledge fail to be advantageous enough to become incorporated, under evolutionary pressure, into the cognitive toolbox of behaving animals? As Virgil's line puts it, *Felix qui potuit rerum cognoscere causas*,³ or, in Dryden's translation:

Happy the man, who, studying nature's laws,
Thro' known effects can trace the secret cause.

Concluding remarks

The thesis advanced by HSP carries with it a lesson for studying the brain basis for vision and other faculties, which the target paper leaves implicit. Namely, instead of focusing exclusively on the function of

³“Fortunate [is he] who was able to know the causes of things” (*Georgics* II:490).

the sensory cortex, which at least in the case of vision can be effectively approximated by unsupervised learning methods (Poggio, Mutch, Leibo, Rosasco, and Tacchetti, 2012; Yamins, Hong, Cadieu, Solomon, Seibert, and Dicarlo, 2014), we should study the entire brain, including sub-cortical circuits that support reinforcement learning (Woergoetter and Porr, 2007; Singh, Lewis, and Barto, 2010; Edelman, 2014).

Finally, it would be interesting to apply HSP's methods to the study of phenomenal awareness, which is the basic component of consciousness. The phenomenological consequence of the HSP thesis is that the world looks the way it does "because evolution." This is fair enough, but in light of the notion that phenomenal awareness is first and foremost the awareness of distinctions and differences (Fekete and Edelman, 2011), and assuming that phenomenality is subject to evolution (Cleeremans, 2008; Ginsburg and Jablonka, 2010), we must hypothesize, in addition, that there are some things about the world that evolving agents get right.

References

- Barlow, H. B. (1980). The absolute efficiency of perceptual decisions. *Proceedings of the Royal Society of London B* 290, 71–82.
- Clark, A. (1993). *Sensory qualities*. Oxford: Clarendon Press.
- Cleeremans, A. (2008). Consciousness: the radical plasticity thesis. In R. Banerjee and B. K. Chakrabarti (Eds.), *Progress in Brain Research*, Volume 168, Chapter 3, pp. 19–33.
- Cosmides, L. and J. Tooby (1997). Evolutionary psychology: A primer. Available online at <http://www.psych.ucsb.edu/research/cep/primer.html>.
- De Pierris, G. and M. Friedman (2013). Kant and Hume on causality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2013 ed.).
- Dobzhansky, T. (1973). Nothing in biology makes sense except in the light of evolution. *The American Biology Teacher* 35, 125–129.

- Edelman, S. (1998). Representation is representation of similarity. *Behavioral and Brain Sciences* 21, 449–498.
- Edelman, S. (1999). *Representation and recognition in vision*. Cambridge, MA: MIT Press.
- Edelman, S. (2008). *Computing the mind: how the mind really works*. New York, NY: Oxford University Press.
- Edelman, S. (2011). The metaphysics of embodiment. *International Journal of Machine Consciousness* 3, 321–325. Part of a collective review of *Embodiment and the Inner Life* by M. Shanahan, Oxford University Press, 2010.
- Edelman, S. (2014, June). Design for a brain? Submitted.
- Edelman, S. and R. Shahbazi (2012). Renewing the respect for similarity. *Frontiers in Computational Neuroscience* 6, 45.
- Fekete, T. and S. Edelman (2011). Towards a computational theory of experience. *Consciousness and Cognition* 20, 807–827.
- Ginsburg, S. and E. Jablonka (2010). Experiencing: a Jamesian approach. *Journal of Consciousness Studies* 17, 102–124.
- Harnad, S. (Ed.) (1987). *Categorical Perception: The Groundwork of Cognition*. New York: Cambridge University Press.
- Hoffman, D. D., M. Singh, and C. Prakash (2014). The interface theory of perception. *Psychonomic Bulletin and Review*. In press.
- Holyoak, K. J. and P. W. Cheng (2011). Causal learning and inference as a rational process: the new synthesis. *Annual Review of Psychology* 62, 135–163.
- Hume, D. (1748). *An Enquiry Concerning Human Understanding*. Available online at <http://eserver.org/18th/hume-enquiry.html>.

- Knill, D. C. and D. Kersten (1991). Ideal perceptual observers for computation, psychophysics and neural networks. In R. Watt (Ed.), *Vision and visual dysfunction*, Volume 14, Chapter 7, pp. 83–97. London: Macmillan.
- Kolodny, O. and S. Edelman (2014, August). The problem of concurrent multimodal serial order in behavior. Submitted.
- Kushnir, T. and A. Gopnik (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science* 16, 678–683.
- Lagnado, D. A. and M. Speekenbrink (2010). The influence of delays in real-time causal learning. *The Open Psychology Journal* 3, 184–195.
- Lagnado, D. A., M. R. Waldmann, Y. Hagmayer, and S. A. Sloman (2007). Beyond covariation: cues to causal structure. In A. Gopnik and L. Schulz (Eds.), *Structure*, pp. 1–48. Oxford University Press.
- Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral Mechanisms in Behavior*, pp. 112–146. New York: Wiley.
- Leising, K. J., J. Wong, M. R. Waldmann, and A. P. Blaisdell (2008). The special status of actions in causal reasoning in rats. *Journal of Experimental Psychology: General* 137, 514–527.
- Macmillan, N. A., H. L. Kaplan, and C. D. Creelman (1977). The psychophysics of categorical perception. *Psychological Review* 84, 452–471.
- Pearl, J. (2001). Causal inference in statistics: A gentle introduction. In *Computing Science and Statistics: Proceedings of Interface '01*, Volume 33. Available online at http://ftp.cs.ucla.edu/pub/stat_ser/R289.pdf.
- Pearl, J. (2009). Causal inference in statistics: an overview. *Statistics Surveys* 3, 96–146.
- Pearl, J. (2013). Structural counterfactuals: A brief introduction. *Cognitive Science* 37, 977–985.
- Poggio, T., J. Mutch, J. Leibo, L. Rosasco, and A. Tacchetti (2012). The computational magic of the ventral stream: sketch of a theory (and why some deep architectures work). CSAIL TR 035, MIT.

- Quine, W. V. O. (1969). Natural kinds. In *Ontological relativity and other essays*, pp. 114–138. New York, NY: Columbia University Press.
- Schaffer, J. (2009). The metaphysics of causation. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2009 ed.).
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science* 237, 1317–1323.
- Shepard, R. N. and S. Chipman (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology* 1, 1–17.
- Singh, S., R. L. Lewis, and A. G. Barto (2010). Intrinsically motivated reinforcement learning: an evolutionary perspective. *IEEE Trans. Auton. Ment. Dev.* 2, 70–82.
- Tenenbaum, J. B. and T. L. Griffiths (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences* 24, 629–641.
- Watanabe, S. (1969). *Knowing and Guessing: A Quantitative Study of Inference and Information*. New York: Wiley.
- Woergoetter, W. and B. Porr (2007). Reinforcement learning. *Scholarpedia* 3(3), 1448.
- Yamins, D. L., H. Hong, C. Cadieu, E. A. Solomon, D. Seibert, and J. J. Dicarlo (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Science* 111, 8619–8624.