

Bayesian inference and “truth”: a comment on Hoffman, Singh, and Prakash

Jacob Feldman

Published online: 18 September 2015
© Psychonomic Society, Inc. 2015

Abstract Hoffman, Singh, and Prakash ([in press](#)) argue that veridicality is neither required nor achieved by the visual system, and propose a new framework in which the literal truth of perceptual inferences plays no role. In this brief comment, I concur with and advocate their basic position, though I go on to argue that Bayesian inference already embodies a similar epistemological stance.

Keywords Bayesian modeling · Visual perception

That “vision usually provides us with a veridical representation of the world” is a cliché so hoary that we vision scientists hardly stop to think about whether it is actually true. Hoffman et al. ([in press](#)) ask us to consider it a bit more carefully. Could such a truism actually be wrong?

It’s worse than wrong—it’s meaningless.

Certainly, we’re all guilty of uttering it, especially in introductory settings. It’s a natural companion to the idea that perception is fundamentally ambiguous, that the proximal stimulus is consistent with infinitely many interpretations. Of course, we routinely go on to add, the visual system usually gives us the *right* interpretation—otherwise, it would seem, we are all hallucinating. Certainly, as Hoffman et al. would agree, the visual system does an exemplary job at resolving the ambiguity. But does it do so by giving us something *true*, or simply something *useful*? Or is this a distinction without a difference?

On that narrow question, it seems to me that Hoffman et al.’s position cannot be disputed: evolution favors fitness, not truth, beauty, or anything else except insofar as it is correlated with fitness. This is literally tautological in the context of Darwinian evolution, as it is essentially a restatement of what is meant by “fitness”—that which is favored by adaptive pressure. So Hoffman et al.’s basic conclusion is inescapable: evolution optimizes fitness, by definition.

The more difficult question is whether true beliefs tend to facilitate fitness. Hoffman et al. give somewhat short shrift to this question, setting up artificial games in which truth and fitness are decorrelated. The result—*inevitably*—is that fitness wins. Truth is irrelevant.

A skeptic might argue that in the real world true percepts and utility-maximizing actions tend to go hand in hand—that being *right* tends to yield tangible rewards. But this is not the case in general. That depends on the utility function, the function that maps decisions to consequences. One might imagine most utility functions place the highest utility on correct inferences, and indeed it is easy to construct payoff matrices where being right always results in higher payoffs than being wrong. Zero/one loss, where you gain utility from each correct classification and lose it from each incorrect classification, is the simplest example.¹ But utility functions in which veridical conclusions automatically convey higher utility are mathematically exceptional, and it is groundless to assume that they predominate in real situations. And again from a Darwinian point of view, when veridicality and fitness are decorrelated, fitness is what matters.

J. Feldman (✉)
Department of Psychology, Center for Cognitive Science,
Rutgers University - New Brunswick, 152 Frelinghuysen Rd,
Piscataway Township, NJ 08854, USA
e-mail: jacob@rucss.rutgers.edu

¹Here I’m assuming that evolutionary fitness is proportional to payoff, which Geisler and Diehl (2002) assume in their simulations. Hoffman et al. distinguish utility from fitness more carefully, but it seems reasonable to assume that they are generally correlated: the more food or money you gain, the better you do in the game of life.

That being said, is it plausible to think that our perceptual conclusions are not veridical? Our intuitions reel at the idea that the world does not actually look the way it appears to. But the conflict with intuition exists *by definition*—of course we find it hard to believe that our *beliefs* are wrong. But our introspections do not count as evidence, no matter how subjectively certain they feel. As scientists, we must ask whether the presumption of veridicality can be defended in more rigorous terms.

In my view, there is no scientific basis for it, and moreover it does not, by itself, really mean anything. To see why, notice that the idea of veridical perception hinges on a presupposition that perceptual judgments per se have truth values—that is, that they can literally match, or fail to match, real-world measurements. But what exactly does this mean?

We have a strong intuition that a statement like “this banana is 20 cm long” is true (or false) in a Platonic sense independent of our measurements. But physicists long ago gave up this idea, in favor of a more rigorous notion of *comparison*. “This banana is 20 cm long” means simply that if we hold the banana up to a 20 cm ruler, they would match in length. It is not independently meaningful to say that the 20 cm ruler is, itself, 20 cm long; it simply acts as a conventional standard.

To assess whether our perceptual beliefs are true, we have to use a similar notion of comparison. Specifically, we have to compare the representation in our heads with the physical state of the world. But a mental representation is a neural state, such as a particular pattern of neural activation; while the physical state of the world is an arrangement of atomic particles. Mental representations and states of the world are, quite concretely, incommensurate. Without substantial additional assumptions, it means nothing to say the mental state “literally” matches the world. The missing element is what philosophers call *semantics*, some stipulation of truth conditions for mental representations (see Feldman, 1999). For example, we could assume that a certain neural state means “the banana is 20 cm long” while another means “the banana is 21 cm long.” But we can’t tell what neural states “mean” just by looking at them. The only evidence that a particular neural state means “the banana is 20 cm long” is that it tends to occur when the observer is looking at a banana that is 20 cm long—coupled with the assumption that the observer is veridical! In other words, the veridicality of the observer’s representation is an *assumption* we have adopted—or, more accurately, a convention of terminology—not an empirical fact. The circularity is inescapable. Because mental states do not have transparent or objectively determinable semantics, it is impossible to say in an independent sense whether they match reality.

What we *can* say, often, is whether estimates from distinct perceptual systems “agree.” When I reach out and

grasp the banana, I may find that it feels the same size as it looked. This is extremely useful, because it means that the perceptual system is *coherent*: estimates drawn from different sorts of evidence agree. But it is not veridicality. Just because many people think that chocolate is delicious does not mean that chocolate is *objectively* delicious; it just means that many people agree with each other. Just because our various sensory modalities tend to converge on mutually consistent perceptual estimates does not mean they are correct—just that they are coherent. Not incidentally, this is the central principle underlying Bayesian inference, which Dennis Lindley (2006) paraphrased as (his emphasis) “BE COHERENT.”

In connection with Bayesian inference, Hoffman et al. go on to argue that conventional Bayesian models of perceptual inference assume that the hypothesis space (the set of models under consideration) is isomorphic to (or to a subset of), the space of possible scenes, from which it would follow that Bayes’ rule involves an estimate of the true state of nature (that is, of which scene is most likely). It is true that many Bayesian perceptual models work this way; in Hoffman et al.’s notation, they assume that $X = W$. But such a strong assumption is not really necessary in a Bayesian framework—at least, it is not required or implied by any of the equations.² Rather, Bayesian inference only assumes that there is some set M of possible models under consideration, which are tied to the data via likelihood functions $p(X|M)$. Bayes’ rule allows these models to be compared *to each other* in terms of plausibility, but says nothing whatsoever about whether any of the models is true in a larger or absolute sense (see Feldman, 2014). The “truth” of the models (whatever that even means—see remarks above about semantics) never enters into it.

This is what is meant by the Bayesian canard, usually attributed to George Box, that “all models are wrong, but some are useful.”³ None of the hypotheses under consideration needs to be *literally* true for the process of Bayesian inference to be useful or coherent, and indeed, in most realistic situations, none are. As Bernardo and Smith (1994) put it: “Nature does not provide us with an exhaustive list of possible mechanisms and a guarantee that one of them is true. Instead, we ourselves choose the list as part of the

²Indeed, in my view, this restrictive interpretation is part of a broader misapprehension of what Bayesian models assume; see Feldman (2013).

³From Box and Tiao (1973) (an influential Bayesian text): “Because we can never be sure that a postulated model is entirely appropriate, we must proceed in such a manner that inadequacies can be taken account of and their implications considered as we go along[.]” (p7). And later (p24): “The only realistic expectation from a statistical analysis is that the conclusions will provide a good enough approximation to the truth.” Good enough, that is, to guide action effectively.

process of settling on a predictive specification that we hope will prove ‘fit for the purpose’[.]” In other words, Bayesian inference does not require—nor, indeed, in any way involve—the literal truth of any hypotheses. All that is needed is that selection of hypotheses guide action “effectively.” And effectiveness, as Hoffman et al. argue, really means fitness. Veridicality is a red herring.

So—at least from a Bayesian perspective—our models of perception do not require that our perceptual impressions of the outside world are usually (or indeed ever) true. So, the skeptic asks, what *does* the world actually look like, if not what it *appears* to look like?

The answer is that it doesn’t “look” like anything. It is a category error to think of the outside world as having a true appearance separable from the interpretations placed upon it by particular observers. We are so immersed in our own subjective interpretation of reality that we confuse it with reality itself (see Koenderink, 2012). In the words of the physicist Arthur Eddington in his influential 1928 monograph. *The Nature of the Physical World*: “I am afraid of this word Reality, not connoting an ordinarily definable characteristic of the things it is applied to but used as though it were some kind of celestial halo. I very much doubt if any one of us has

the faintest idea of what is meant by the reality or existence of anything but our own Egos.”

References

- Bernardo, J.M., & Smith, A.F.M. (1994). *Bayesian theory*. Chichester: John Wiley.
- Box, G.E.P., & Tiao, G.C. (1973). *Bayesian inference in statistical analysis*. Reading, Massachusetts: Addison-Wesley.
- Eddington, A.S. (1928). *The nature of the physical world*. New York: Macmillan.
- Feldman, J. (1999). Does vision work? Towards a semantics of perception. In Lepore, E., & Pylyshyn, Z. (Eds.) *What is Cognitive Science?* (pp. 208–229). New York: Basil Blackwell.
- Feldman, J. (2013). Tuning your priors to the world. *Topics in Cognitive Science*, 5(1), 13–34.
- Feldman, J. (2014). Bayesian models of perceptual organization. In Wagemans, J. (Ed.) *Handbook of perceptual organization*. Oxford: Oxford University Press.
- Geisler, W.S., & Diehl, R.L. (2002). *Bayesian natural selection and the evolution of perceptual systems* (Vol. 357, pp. 419–448).
- Hoffman, D.D., Singh, M., Prakash, C. (in press). The interface theory of perception. *Psychonomic Bulletin and Review*.
- Koenderink, J. (2012). *Visual awareness*: Clootrans Press.
- Lindley, D.V. (2006). *Understanding uncertainty*. Hoboken: John Wiley.