Chapter 2: What do we mean by "The structure of the world"?

DONALD D. HOFFMAN

The first sentence of Jepson, Richards, and Knill's insightful paper reads as follows: "The world we live in is a very structured place." I wish to meditate on this sentence.

What do we mean when we speak of the "structure of the world?" In particular, what do we mean when we, as purveyors of a Bayesian approach to perception, speak of the "structure of the world?"

There is, of course, a strictly mathematical answer to this question. We can say that the structure of the world is a measure, which Bayesians call the prior, on a measurable space. It is through prior measures that the structure of the world is expressed in computations using Bayes formula.

This is surely true and interesting, but misses the real point of my question. Given this mathematical answer, the question is how shall we conceive of the world represented by a prior? Is it a world independent of the observer? Is it a world whose structure we can assess objectively and then compare, favorably or not, with the observer's perceptions? Or is it an observer-dependent world?

I think most of us would vote for a world independent of the observer. We admit this in the terminology we use. We speak of perception as "generating estimates of world properties" (Knill & Kersten, 1995) or as "recovering world properties," as though world properties are objectively out there, independent of the observer, and the task of the observer is to match its perceptions, as best it can, to these properties. Marr speaks of the senses as providing "perception of the real world outside" (Marr, 1982, p. 29). The world is out there, whether we look or not, and the goal of perception is to estimate its structure. This is a view I too have espoused, asserting once that a central problem of perception is how it remains "true to the real world" (Hoffman, 1983, p. 154).

But I now think this is mistaken, and Bayes tells us why. Perception is probabilistic inference. Observers acquire probabilistic premises and reach probabilistic conclusions. What an observer sees, and all it can ever see, are its own probabilistic conclusions. When I look at a table, the 3D shape I see is the conclusion of inferences involving stereo, motion, shading, and texture. The color I see is my conclusion. The temperature, hardness, and smoothness I feel when I touch it are conclusions of my somatosensory inferences. The sound I hear when I tap it is a conclusion of my auditory inferences. In short, the table, and all properties of it that I experience, are my conclusions. What holds for tables holds also for forks, suns, brains, and

neurons. These are the products of perception, not the antecedents. In perception, as a Bayesian would put it, we perceive only our posteriors.

This is, of course, not a new idea. But it is sometimes difficult to swallow. As Crick puts it in his *Astonishing Hypothesis* (1994, p. 33), "It is difficult for many people to accept that what they see is a symbolic interpretation of the world – it all seems so like 'the real thing.' But in fact we have no direct knowledge of objects in the world."

So when, as Bayesians, we examine the "external world" to determine what priors we should use, what do we find? We find our own posteriors. And nothing else. All we can ever see in perception is our own posteriors.

Is there nevertheless an observer-independent world out there? I think so. My perceptions are so systematic (I can use group theory to predict what I'll see if I move my head) and intransigent (I can't walk through walls) that I suspect they are due, in part, to something independent of me.

But does this observer-independent world resemble what I see, hear, feel, or smell? That is more than I can know. But I suspect it does not. We all suspect it does not in the case of synesthetics. When we hear of someone who, upon tasting mint, feels as though he were grasping with his hands tall, smooth, cool, columns of glass (Cytowic, 1993) we suspect that there is no relation of resemblance between his perceptions and the observer-independent world with which he might be interacting. But why should we think that the taste of mint that *we* perceive is any more likely to resemble that observer-independent world? Russell (1912, p. 33), when considering this kind of question, argued that we could at least assume that *depth order* (say front to back) as we perceive it, resembles the true order in the observer-independent realm. But he had apparently never seen a Necker cube. The perceived order of its faces changes as we switch from one perceptual conclusion to another. We have no reason to suppose there is a concomitant change in order of an observer-independent realm. And we have no reason to suppose, in any case of perception, that what we perceive bears any resemblance to the observer-independent world.

In sum, for good Bayesian analysis we need appropriate priors (and likelihoods). But when we look we see only our posteriors. What are we to do? Well, what we in fact do is to fabricate those priors (and likelihoods) which best square with our posteriors. And we are happy when the three are finally consistent, via Bayes. Jepson, Richards, and Knill's paper is an excellent example of how to do this.

But if we think that what we are really doing is getting (perhaps through high-tech physical devices) the true properties of the observer-independent world and from these deducing the appropriate priors and likelihoods, we should think again.

References

Crick, F.H.C. (1994). *The Astonishing Hypothesis.* New York: Macmillan.
Cytowic, R.E. (1993). *The Man Who Tasted Shapes.* New York: Putnam.
Hoffman, D.D. (1983). The interpretation of visual illusions. *Scientific American*, **249**, 154-162.
Knill, D.C. & Kersten, D. (1995). Implications of a Bayesian formulation of visual information processing for psychophysics. In *Perception as Bayesian inference*, ed. W. Richards and D. Knill. Cambridge: Cambridge University Press.
Marr, D. (1982). *Vision.* San Francisco: Freeman.
Russell, B. (1912). *The Problems of Philosophy.* Oxford: Oxford University Press.