

The New York Times

Opinionator

JANUARY 15, 2012, 5:00 PM

Do Thrifty Brains Make Better Minds?

By **ANDY CLARK**

[The Stone](#) is a forum for contemporary philosophers on issues both timely and timeless.

Tags:

[Neuroscience](#), [perception](#), [Philosophy](#), [vision](#)

Might the miserly use of neural resources be one of the essential keys to understanding how brains make sense of the world? Some recent work in computational and cognitive neuroscience suggests that it is indeed the frugal use of our native neural capacity (the inventive use of restricted “neural bandwidth,” if you will) that explains how brains like ours so elegantly make sense of noisy and ambiguous sensory input. That same story suggests, intriguingly, that perception, understanding and imagination, which we might intuitively consider to be three distinct chunks of our mental machinery, are inextricably tied together as simultaneous results of a single underlying strategy known as “predictive coding.” This strategy saves on bandwidth using (who would have guessed it?) one of the many technical wheezes that enable us to economically store and transmit pictures, sounds and videos using formats such as JPEG and MP3.

In the case of a picture (a black and white photo of Laurence Olivier playing Hamlet, to activate a concrete image in your mind) predictive coding works by assuming that the value of each pixel is well predicted by the value of its various neighbors. When that’s true — which is rather often, as gray-scale gradients are pretty smooth for large parts of most images — there is simply no need to transmit the value of that pixel. All that the photo-frugal need transmit are the deviations from what was thus predicted. The simplest prediction would be that neighboring pixels all share the same value (the same gray scale value, for example), but much more complex predictions are also possible. As long as there is detectable regularity, prediction (and hence this particular form of data compression) is possible.

Leif Parsons

Such compression by informed prediction (as Bell Telephone Labs first discovered back in the 1950s) can save enormously on bandwidth, allowing quite modest encodings to be reconstructed, by in effect “adding back in” the successfully predicted elements, into rich and florid renditions of the original sights and sounds. The basic trick is one we can use in

daily life, too. Suppose you make a plan with your friend Duke by saying that if you *don't* call him, then all is “as expected” and that he should therefore meet your plane at Miami airport next Wednesday at 9 a.m. local time. Your *failure* to call is then (technically speaking) a tiny little one-bit signal that conveys a large amount of neatly compressed information! The trick is trading intelligence and foreknowledge (expectations, informed predictions) on the part of the receiver against the costs of encoding and transmission on the day.

A version of this same trick may be helping animals like us to sense and understand the world by allowing us to use what we already know to predict as much of the current sensory data as possible. When you think you see or hear your beloved cat or dog when the door or wind makes just the right jiggle or rustle, you are probably using well-trained prediction to fill in the gaps, saving on bandwidth and (usually) knowing your world better as a result.

Neural versions of this predictive coding trick benefit, however, from an important added dimension: the use of a stacked hierarchy of processing stages. In biological brains, the prediction-based strategy unfolds within multiple layers, each of which deploys its own specialized knowledge and resources to try to predict the states of the level below it.

This is not easy to imagine, but it rewards the effort. A familiar, but still useful, analogy is with the way problems and issues are passed up the chain of command in rather traditional management hierarchies. Each person in the chain must learn to distil important (hence usually surprising or unpredicted) information from those lower down the chain. And they must do so in a way that is sufficiently sensitive to the needs (hence, expectations) of those immediately above them.

In this kind of multilevel chain, all that flows upward is *news*. What flows forward, in true bandwidth-miser style, are the deviations (be they for good or for ill) from each level's predicted events and unfoldings. This is efficient. Valuable bandwidth is not used sending well-predicted stuff forward. Why bother? We were expecting all that stuff anyway. Who in corporate headquarters wants to know that the work of Jill/Jack proceeded exactly as expected? Instead, that expensive bandwidth is used only to flag what may more plausibly demand attention: outcomes that gloriously exceeded or sadly fell short of expectations. Things work similarly — if the predictive coding account is correct — in the neural incarnation. What is marked and passed forward in the brain's flow of processing are the divergences from predicted states: divergences that may be used to demand more information at those very specific points, or to guide remedial action.

All this, if true, has much more than merely engineering significance. For it suggests that perception may best be seen as what has sometimes been described as a process of “controlled hallucination” (Ramesh Jain) in which we (or rather, various parts of our brains) try to predict what is out there, using the incoming signal more as a means of tuning

and nuancing the predictions rather than as a rich (and bandwidth-costly) encoding of the state of the world. This in turn underlines the surprising extent to which the structure of our expectations (both conscious and non-conscious) may quite literally be determining much of what we see, hear and feel.

The basic effect hereabouts is neatly illustrated by a simple but striking demonstration (used by the neuroscientist Richard Gregory back in the 1970's to make this very point) known as "the hollow face illusion." This is a well-known illusion in which an ordinary face mask viewed from the back can appear strikingly convex. That is, it looks (from the back) to be shaped like a real face, with the nose sticking outward rather than having a concave nose cavity. Just about any hollow face mask will produce some version of this powerful illusion, and there are many examples on the Web, like this one:

The hollow face illusion illustrates the power of what cognitive psychologists call "top-down" (essentially, knowledge-driven) influences on perception. Our statistically salient experience with endless hordes of convex faces in daily life installs a deep expectation of convexness: an expectation that here trumps the many other visual cues that ought to be telling us that what we are seeing is a concave mask.

You might reasonably suspect that the hollow-face illusion, though striking, is really just some kind of psychological oddity. And to be sure, our expectations concerning the convexity of faces seem especially strong and potent. But if the predictive-coding approaches I mentioned earlier are on track, this strategy might actually pervade human perception. Brains like ours may be constantly trying to use what they already know so as to predict the current sensory signal, using the incoming signal to constrain those predictions, and sometimes using the expectations to "trump" certain aspects of the incoming sensory signal itself. (Such trumping makes adaptive sense, as the capacity to use what you know to outweigh some of what the incoming signal seems to be saying can be hugely beneficial when the sensory data is noisy, ambiguous, or incomplete — situations that are, in fact, pretty much the norm in daily life.

This image of the brain (or more accurately, of sensory and motor cortex) as an engine of prediction is a simple and quite elegant one that can be found in various forms in contemporary neuroscience. For useful surveys, see Kveraga et al (2007), Bubic et al (2010), and for my own favorite incarnation, see Friston (2010). It has also been shown, at least in restricted domains, to be computationally sound and practically viable. Just suppose (if only for the sake of argument) that it is on track, and that perception is indeed a process in which incoming sensory data is constantly matched with "top down" predictions based on unconscious expectations of how that sensory data should be. This would have important implications for how we should think about minds like ours.

First, consider the unconscious expectations themselves. They derive mostly from the

statistical shape of the world as we have experienced it in the past. We see the world by applying the expectations generated by the statistical lens of our own past experience, and not (mostly) by applying the more delicately rose-nuanced lenses of our political and social aspirations. So if the world that tunes those expectations is sexist or racist, future perceptions will also be similarly sculpted — a royal recipe for tainted evidence and self-fulfilling negative prophecies. That means we should probably be very careful about the shape of the worlds to which we expose ourselves, and our children.

Second, consider that perception (at least of this stripe) now looks to be deeply linked to something not unlike imagination. For insofar as a creature can indeed predict its own sensory inputs from the “top down,” such a creature is well positioned to engage in familiar (though perhaps otherwise deeply puzzling) activities like dreaming and some kind of free-floating imagining. These would occur when the constraining sensory input is switched off, by closing down the sensors, leaving the system free to be driven purely from the top down. We should not suppose that all creatures deploying this strategy can engage in the kinds of self-conscious deliberate imagining that we do. Self-conscious deliberate imagining may well require substantial additional innovations, like the use of language as a means of self-cuing. But where we find perception working in this way, we may expect an interior mental life of a fairly rich stripe, replete with dreams and free-floating episodes of mental imagery.

Finally, perception and understanding would also be revealed as close cousins. For to perceive the world in this way is to deploy knowledge not just about how the sensory signal should be right now, but about how it will probably change and evolve over time. For it is only by means of such longer-term and larger-scale knowledge that we can robustly match the incoming signal, moment to moment, with apt expectations (predictions). To know that (to know how the present sensory signal is likely to change and evolve over time) just *is* to understand a lot about how the world is, and the kinds of entity and event that populate it. Creatures deploying this strategy, when they see the grass twitch in just that certain way, are already expecting to see the tasty prey emerge, and already expecting to feel the sensations of their own muscles tensing to pounce. But an animal, or machine, that has *that kind of grip* on its world is already deep into the business of understanding that world.

I find the unity here intriguing. Perhaps we humans, and a great many other organisms, too, are deploying a fundamental, thrifty, prediction-based strategy that husbands neural resources and (as a direct result) delivers perceiving, understanding and imagining in a single package?

Now there's a deal!

References

Bubic A, von Cramon DY and Schubotz RI (2010) Prediction, cognition and the brain. *Front. Hum. Neurosci.* 4:25: 1-15

Friston K. (2010) The free-energy principle: a unified brain theory? *Nature Reviews: Neuroscience* 11(2):127-38.

Helmholtz, H. (1860/1962). *Handbuch der physiologischen optik* (Southall, J. P. C. (Ed.), English trans.), Vol. 3. New York: Dover.

Kveraga, K., Ghuman, A.S., and Bar. M. (2007) Top-down predictions in the cognitive brain. *Brain and Cognition*, 65, 145-168

Andy Clark is professor of logic and metaphysics in the School of Philosophy, Psychology and Language Sciences at Edinburgh University, Scotland. He is the author of [Supersizing the Mind: Embodiment, Action, and Cognitive Extension](#)” (Oxford University Press, 2008).

Copyright 2012 The New York Times Company | [Privacy Policy](#) | [NYTimes.com](#) 620 Eighth Avenue New York, NY 10018