AVANT, Vol. XI, No. 2
ISSN: 2082-6710 avant.edu.pl/en
DOI: 10.26913/avant.2020.02.15

Comment on "Thinking about Semantic Information"

Daniel C. Dennett
Center for Cognitive Studies,
Tufts University, Medford, MA USA

Received 30 June 2020; accepted 9 July 2020; published 22 July 2020

The commented paper: Miłkowski, M. (2020). Thinking about Semantic Information. *Avant*, 11(2). doi: 10.26913/avant.2020.02.08

Keywords: semantic information; information theory; encoding; satisfaction conditions; information measures; From Bacteria to Bach and Back

Marcin Milkowski provides an insightful overview of work on information theories over the years, from Gabor and Shannon to Wolpert and Kolchinski. He accurately summarizes my work, distilling four claims that I made in *BBB,* and argues that my arguments in favor of the last two of these are invalid.

> 3. The amount of semantic information carried or contained in any delimited episode or item is also not usefully measurable in units but roughly comparable in local circumstances.
>
> 4. Semantic information need not be encoded to be transmitted or saved (Dennett, 2017, p. 128).

Milkowski says "In contrast, I claim that both measurement and encoding are interrelated ideas that are indispensable in any serious, biologically plausible account of cognitive processing." (ms p1) But I agree with this. Milkowski does a fine job outlining the roles of measurement and encoding that must be played in any theory of cognition, but it doesn't follow from these points that *semantic* information is either measurable or coded. Consider an old example of mine: "a wagon with spoked wheels carries not only grain or freight from place to place; it also carries the brilliant idea of a wagon with spoked wheels." (*Darwin's Dangerous Idea*, p. 348) The details will bring out the point. Let's say several thousand years ago one evening a wagon with spoked wheels rolled through a little town in Asia minor, a novelty never seen before. By morning it had long disappeared and, as it happened, only four pairs of retinas were briefly stimulated by the reflected light from the wheels: one dog, one donkey, one slightly drunk fisherman and one clever thirteen-year-old boy. In each case, retinal stimuli were processed in the normal way—and I'll stipulate for the sake of the example that these processes were both coded and measurable, in bits or logons or whatever you like. But in only one brain

did the *idea* of a spoked wheel register, somewhat dimly, but with enough of the zing of novelty to inspire the lad to ponder and reflect and sketch and eventually fashion a spoked wheel. The idea spread and the little town became the source of spoked wheels in the area. What triggered this readily detectable economic and technological explosion? Information. As Gibson would say, the information was in the light, but the dog's eyes were no more able to *use* the information than the "evil eyes" painted on a village wall, on which the light with the information also fell. How much information (measured how) did the boy acquire? What was the code? It is not clear that we should try to answer them. Whatever measure we used would be largely irrelevant to the effect on the spread of the meme for spoked wheels, and the codes used by different individuals as they picked up the meme would no doubt be different.

The reason why semantic information is so elusive to measurement and coding is that we human beings have turned ourselves into super-detectors with huge stores in onboard *information* ready to use in interpreting what our eyes and ears (mainly) detect. Language has a lot to do with this, of course. Without language, which undergirds cumulative cultural evolution, individual brains would be deeply ignorant, maybe more knowing than a dog or a donkey but not orders of magnitude more knowing. We know so much that we can set up what might be called "nonce-codes" at will, and without even prior arrangement (which Shannon's measure requires in its original articulation). You want to know who won the World Cup final game and I step into the bar and say "No samba tonight!" or "Happy Rio!" (or "Arc de Triomphe" or "the piranhas are gorged" or . . . .) and you know that Brazil either lost to France or vice versa. (I could achieve the same communication to a well-informed audience by triumphantly brandishing a *pao de queijo* or *baguette*.) A contrived example, like my Trafalgar Square case, but it is extreme in order to make the point that so much of what we learn depends on what information we already have, and especially on what we share. There are untold thousands of ways I can inform you of some particular fact, and thousands of ways you can confirm to me that you have "got it". A teacher can say "the sum of the squares of the sides is equal to the square of the hypotenuse" to a group of young students and know that only a fraction of them will acquire the information. Nods and frowns and non sequiturs are signals—are they all *encoded?* How many codes do we know? Billions?

Switching to other species, consider the fact that birds that have never seen a nest being built will build a species-typical nest on first attempt. The information to build that nest is somehow embodied in the bird, but encoded? Maybe, if we stretch the notion of coding and decoding into some intricate convolution, but the assumption that the quest to discover how this is accomplished is looking to *find and decipher a code* for building that nest is beyond optimistic. Note that this is a nice case to evaluate with both Floridi's and MacKay's definitions: there is "instruction" (Floridi) that the bird doesn't need to get elsewhere because it already has the knowhow, and it plays a "selective function on the range of states of your conditional readiness" (MacKay) in that the bird builds a specific, readily identifiable nest. Yes, there must be coding of various sorts in the birds' nervous systems, and yes, there must be tweaks or biases or other influencers that modulate the processing of those codes, but to call those tweaks a code is, again, to multiply codes beyond necessity. There is also obviously some kind of "search image" that enables animals to distinguish potential mates from members of other species but is this image a code? There are orchids that "fool" wasps or bees into "pseudocopulation," and

this clearly involves *misinforming* the pollinators. Do we want to say that the orchid doesn't know it, but it has "decoded" the "code" of the pollinator? Perhaps, but I'm not convinced.

I'm glad that Milkowski discusses my Trafalgar Square intuition pump, which I think is sometimes misinterpreted. He notes that by my own earlier account in *Content and Consciousness* I should say that their information is only *similar*. But I do note just that, in part of the quotation Milkowski provides: " . . . *that proposition* did not, we can suppose, 'occur to' any of them, and even if it had, *it would have had very different import* [italics added] for Jacques, Sherlock, Tom, and Boris." And yet still, I insisted, they shared some semantic information. How did I establish this? By imagining a quiz show, in three different languages (English, French, Russian) in which they are to respond to an expression of the proposition that a Frenchman has murdered somebody in Trafalgar Square by pushing the affirmative or negative button. Contrived and extreme. But that is what language permits us to do! I could generate any number of other little scenarios where precisely this believed proposition explains why all four of them act the same. Language creates (at least in idealization) *propositions* that can be indefinitely fine-grained and specific, and that can be expressed in different languages that are reliably translatable . My "Two Black Boxes" intuition pump in *Darwin's Dangerous Idea* exploits this to describe a mysterious regularity in the activity of two computers that *do not share the same code* but instead "use the world as a 'one-time pad'" as Danny Hillis pointed out. (p. 421). That thought experiment demonstrates the ineliminable reliance on *aboutness* in explaining some (contrived) regularities in the world, an extreme case of something everywhere observable in less carefully highlighted phenomena. On the same page I give another example of a code that two earthlings could use to talk in the presence of "Martians" who had Laplacean knowledge of their brains, but without knowing the world they inhabit, they would be unable to decode the messages. Milkowski writes:

> For formal systems, we surely understand what correspondence theory of truth may be like. The problem is how to go from formal systems to sub-personal representations, conscious thought, and natural language. The promise of naturalistic accounts of semantic information is that they could supply the solution to this problem.

I think this optimism is misguided. Milkowski is right that we have a good handle on correspondence theories of truth for formal systems, but that is because they are artificial, top-down-designed sets of conventions that can be made as hard-edged as you please. In such systems, meanings can be *stipulated.* Nature doesn't work that way. Nature designs and builds mechanical systems that reliably track affordances and lets evolution prune and shape these mechanisms opportunistically. Not only are the reference boundaries between "terms" fuzzy; the distinction between a code and a simpler, uncoded detection device is fuzzy. Gradualism is everywhere in the natural world, and the bright lines we draw for convenience are usually quite arbitrary. The digitization that has emerged with natural languages—the correction-to-norms of phonemes, the local speciation events (Richard, 2020) that divide the meanings of once-shared terms in two, with one group adopting one meaning and another group adopting the other meaning—has created the persistent illusion that we can rigorously define *concepts* and *propositions,* and have a crystalline theory of meaning. In some abstract and artifactual topics like geometry and arithmetic we can achieve this; no integer is only *sorta* odd or even, and no geometric line is *sorta* straight, so we can generate ironclad proofs and derivations that

then provide us with an ideal model, but the hope of extending this to the semantic information exploited by time-pressured, fidgety, scrambling animals in the rough-and-tumble physical world is, I think, quixotic.

References

Dennett, D.C. (1969). *Content and Consciousness*. London and New York: Routledge Classics.

Dennett, D. C. (1995). *Darwin's Dangerous Idea*. New York: Simon & Schuster.

Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. New York: WW Norton & Company.

Floridi, L. (2010). *Information: A very short introduction*. Oxford: Oxford University Press.

MacKay, D. M. (1969). *Information, mechanism and meaning*. Cambridge: M.I.T. Press.

Miłkowski, M. (2019). Thinking about Semantic Information. *Avant*, 11(2).

Richard, M. (2020). *Meanings as Species*. Harvard: Harvard University Press.

The editorial and publishing process of this publication has been financed by the Ministry of Science and Higher Education from the funds for the dissemination of research (DUN) within the framework of publishing activity, contract no. 711/P-DUN/2019, period of implementation: the years 2019–2020.