



Thinking about Semantic Information

Marcin Milkowski 

Institute of Philosophy and Sociology
Polish Academy of Sciences
mmilkows@ifispan.waw.pl

Received 25 February 2019; accepted 4 April 2019; published 23 February 2020

Abstract

In his recent book, Daniel Dennett defends a novel account of semantic information in terms of design worth getting (Dennett, 2017). While this is an interesting proposal in itself, my purpose in this commentary is to challenge several of Dennett's claims. First, he argues that semantic information can be transferred without encoding and storing it. Second, this lack of encoding is what makes semantic information unmeasurable. However, the argument for both these claims, presented by Dennett as an intuition pump, is invalid.

Keywords: semantic information; information theory; encoding; satisfaction conditions; information measures

1. Introduction

Dennett's recent book offers a fascinating new vision of a biologically evolved cognition. Cognition is no longer conceived of as individual achievement of singular minds, but in terms of forced moves in the design space of biological evolution. As Dennett has stressed over the years, rationality need not be thought of as a property of individuals, but can be considered in terms of free-floating rationales (Dennett, 1995). My focus here will be on a particular issue in this grand scheme: semantic information. First, I will briefly sketch the background of Dennett's proposal for understanding semantic information in terms of design worth getting. In particular, this proposal includes claims that design worth getting is not measurable and need not be encoded in order to be transmitted or saved. First, I summarize the proposal to understand semantic information in terms of design worth getting. In the following section, my focus will be on an intuition pump proposed by Dennett, which is supposed to substantiate these claims, but fails to do so. Then, I briefly point out why the idea of encoding has become notorious in the study of cognition. In contrast, I claim that both measurement and encoding are interrelated ideas that are indispensable in any serious, biologically plausible account of cognitive processing.

2. Semantic Information as Design Worth Getting

To make any account of cognition at least faintly cogent, one should understand how cognition could aim at truth. This is virtually impossible without understanding what semantic information is.

The proposal defended in the book is that semantic information is design worth getting (Dennett, 2017, Chapter 6). Dennett defends the following claims about semantic information:

1. Semantic information is valuable—misinformation and disinformation are either pathologies or parasitic perversions of the default cases.
2. The value of semantic information is receiver-relative and not *measurable* in any non-arbitrary way but can be *confirmed* by empirical testing.
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).

The issue of whether semantic information is inherently valuable is interesting in its own regard but I will set it aside for the purposes of the current discussion. I will deny claims 3 and 4 above because the argument that is supposed to justify them is invalid. Moreover, it seems *prima facie* impossible to defend. Let me elucidate.

One of the tasks of a theory of semantic information is to specify what it takes for information to be *about* something. To understand this task, it is helpful to contrast it with what the mathematical theory of information is supposed to provide. Shannon stressed that semantic aspects of information are irrelevant for his theory (Shannon, 1948). At the same time, the fact that the purpose of Shannon's theory is to solve the fundamental communication problem—that of “reproducing at one point either exactly or approximately a message selected at another point”—does not imply that Shannon's information does not carry any information about anything. Obviously, if it did not carry information about anything, nobody would be in the business of communicating it. The question is whether Shannon's theory of information is able to elucidate semantic aspects of information.

To answer this question, one must precisely understand what is meant by ‘semantic aspects of information’. Floridi's introductory text provides the following answer: “When data are well formed and meaningful, the result is also known as semantic content” (Floridi, 2010, p. 34). While Floridi subsequently distinguishes two kinds of semantic information, instructional and factual, the notion of meaning remains somewhat unclear. But it seems that Floridi's account relies on the idea of satisfaction conditions: namely, information has semantic contents if and only if it can be satisfied, either as instruction (followed by some system) or as factual claim (thus, truth-evaluable). Thus understood, semantic information is meaningful, or satisfiable information. The task of a theory that elucidates this notion is to specify what makes a vehicle of information to have satisfaction conditions.

Some analyze semantic information in an action-oriented way rather than in terms of satisfaction conditions. For example, meaning was analyzed by Donald M. MacKay, one of Dennett's heroes of the nearly forgotten cybernetic past, in terms of dispositions to act: “meaning of a message can be defined very simply as its selective function on the range of the recipient's

states of conditional readiness for goal-directed activity; so that the meaning of a message to you is its selective function on the range of your states of conditional readiness” (MacKay, 1969, p. 24). Thus, whenever a physical system responds to a state of a vehicle of information by changing its readiness to act, which is reducible to changing its conditional probabilities of further action, then this vehicle of information is meaningful. Interestingly, this means that it takes very little to have semantic information in this sense. Any communication channel in Shannon’s sense, whenever it includes a receiver that is capable of anything in response to messages it receives, contains meaningful information in MacKay’s sense. The bar for semantic information is quite low, it seems, as this kind of information does not involve truth or satisfaction conditions, at least not in any straightforward way. This idea seems to recur in contemporary attempts to bring together information and autonomous systems: as long as a dynamical system is in the business of maintaining its existence, the information it receives could be thought to be semantic if it contributes to this task (Kolchinsky & Wolpert, 2018).

Dennett’s proposal seems to be clearly in this latter, action-oriented tradition. Instead of focusing on how information could have satisfaction conditions, he is concerned with the question of how it is used and what makes it valuable. But the latter question is unlikely to be complete without actually providing a theory of aboutness of information (in terms of its satisfaction conditions or otherwise), however daunting this task might be. Regardless, the crux of the problem of defending an account of semantic information is that it must start somewhere. Given that information is already well-understood in non-semantic terms, the general agreement is that an account of semantic information must start with non-semantic information.

And herein lies trouble for Dennett. One cannot speak of non-semantic information without speaking of encoding. One implies the other. There is no reasonable or workable notion of non-semantic information that is measure-free. The whole purpose of devising the mathematical theory of communication was, after all, to measure information. So, at least *prima facie*, it seems that Dennett wants to claim that semantic aspects of information somehow make it no longer measurable, at least not in its semantic part, and that it is encoding-free.

To reiterate, the problem of devising a theory of semantic information is to start with a respectable theory of non-semantic information and provide a clear account of what makes it semantic. But respectable theories of non-semantic information are based on the idea of encoding, and they provide ways to measure information. This is true of Shannon’s notion of information, but it is equally true of an even more primitive notion: structural information (or *logon* information, to use Gabor’s (1946) terminology), which can be specified without the notions of a sender, channel, receiver and probability distributions of messages, which are required for Shannon’s theory (but note that both are frequently confused, in particular by philosophers; cf. Harms (1998)). Arguably, Gabor’s notion constitutes the minimal notion of information. Its measure could be defined in several equivalent ways, as logical dimensionality of vehicles of any form, minimum equivalent number of independent features which must be specified, or simply the number of degrees of freedom of a vehicle (MacKay, 1969). The logical dimensionality of X is the number of yes/no questions one must ask to fully describe X. It is, hopefully, clear that structural information is defined in a way that makes it measurable: we have to establish the number of degrees of freedom of a vehicle, or establish the total number of questions that must be asked in order to describe the vehicle’s possible states. Moreover,

the level of degrees of freedom is equated here with the number of symbols in the physical encoding. No wonder that semantic theories of information usually assume that it is encoded and measurable. Thus, Dennett owes us an argument to show why his notion of semantic information is both independent of encoding and non-measurable.

3. Jacques, Sherlock, Tom, and Boris share the same truth, or do they?

To establish that semantic information does not require encoding and is not apt for measurement, Dennett presents the following intuition pump:

Jacques shoots his uncle dead in Trafalgar Square and is apprehended on the spot by Sherlock. Tom reads about it in the *Guardian* and Boris learns of it in *Pravda*. Now Jacques, Sherlock, Tom, and Boris have had remarkably different experiences, but there is one thing they share: semantic information to the effect that a Frenchman has committed a murder in Trafalgar Square. They did not all *say* this, not even “to themselves”; *that proposition* did not, we can suppose, “occur to” any of them, and even if it had, it would have had very different import for Jacques, Sherlock, Tom, and Boris. They share no encoding, but they do share semantic information (Dennett, 2017, pp. 112–113).

This story is supposed to establish the following claim: that Jacques, Sherlock, Tom, and Boris share the same piece of semantic information without sharing the same encoding. Interestingly, in his first book, *Content and Consciousness*, Dennett bluntly denied that multiple agents could easily share the *same* semantic information, and relied on MacKay’s account of meaning, which is receiver-relative, not absolute (Dennett, 1969, Chapter X). Even the example of cross-linguistic communication is similar:

Both speaker and hearer must share relevant knowledge for communication to occur. There is more to this than the fact that if I attempt to communicate in English with a person who speaks no English, I will not succeed in producing in him information similar to mine at all. Even if the hearer is English, he must also have much the same background of information on the subject of discussion as I have (Dennett, 1969, p. 188).

So, according to Dennett’s claims in *Content and Consciousness*, Jacques, Sherlock, Tom, and Boris have only *similar* information, but they clearly have diverging background information. Moreover, the effects of information most likely differ for all four. Jacques, if psychopathic, could enjoy killing his uncle and cry with joy when thinking about it. Sherlock apprehends him when he sees Jacques red-handed. Tom may think that Brexit should happen to stop these darn Frenchmen from coming to Trafalgar Square, and Boris could marvel about the demise of Western civilization that only the Communist Party could salvage. For an action-based definition of meaning, such as MacKay’s, the same event, perceived and reported in various languages, has quite different semantic effects (which is what Dennett seems to imply in his story).

In his new account, to claim that they still share the same information, Dennett points out that there is something that they all share: “that a Frenchman has committed a murder in Trafalgar Square” (note that this is not necessarily contradictory with his previous account: after all, similarity of semantic information may involve some common features). Let’s suppose indeed

that Sherlock knows that Jacques is French (Tom and Boris may guess it from his name). The best way to elucidate Dennett's claim would be to say that they could entertain a thought whose conditions of satisfaction are such that a Frenchman has committed a murder in Trafalgar Square. The problem is that Dennett's account of semantic information does not provide the account of satisfaction conditions, so the point of the intuition pump is difficult to justify using merely his own take on semantic information. But let's suppose we could appeal to some other account of satisfaction conditions, in terms of indication, local natural information, or structural similarity, for example (Milkowski, 2015).

These four people could share the same piece of information with the same satisfaction conditions without sharing the same encoding, as long as differing pieces of information simply have the same satisfaction conditions. This is very easy to achieve. For example, all logical truths have the same satisfaction conditions (as they are always true), although their vehicles are different. And it's easy to point out that there are multiple ways of describing any state of affairs, which means that this point is not limited to logical truths, but also covers contingent statements. So, certainly, Tom and Boris could have the same semantic information and differing encoding thereof. The same proposition need not be shared. What is important is that the satisfaction conditions of some of their mental representations be the same.

But this means that the intuition pump hardly justifies Dennett's claim that semantic information "need not be encoded to be transmitted or saved". It only implies that differing pieces of information can be semantically satisfied in the same conditions, while being formed in differing ways, or encoded in differing ways. As Floridi makes clear in his definition, semantic information is a kind of *well-formed* data. In other words, semantic information cannot float free of its encoding. Dennett's intuition pump is not enough to show that it could.

We might be tempted to think that all there is to sharing semantic information is simply some kind of intentional ascription of beliefs to intentional systems. And since beliefs are merely ascribed in the intentional stance (Dennett, 1987), they need not have any encoding in Jacques, Sherlock and so on. Voilà! There is no straightforward connection between intentionality and internal cognitive states in Dennett's account, right?

Not so quick. The explanatory buck does not stop at intentionality. To causally explain why all four could have the same belief—that a Frenchman killed his uncle in Trafalgar Square—we would have to focus on the cognitive processing that they were involved in. Dennett's account of intentional stance does not deny this at all. It does not obviate the need to use the design stance in order to elucidate functionality, or the physical stance to explain physical facts.

Obviously, the amount of cognitive processing caused by linguistic utterances (in Tom's and Boris's case), action-monitoring (in Jacques), or perception (in Sherlock's case) is difficult to establish in an ecologically valid fashion right now. But we should not be deluded into thinking that it does not require processing physical vehicles of information by the fact that a fair amount of linguistic processing happens without our explicit knowing. Thus, it could also be the case that inferential processes caused by linguistic utterances lead Tom and Boris to infer that Jacques is French, for example. It is very difficult to see what one could mean that they

understand an utterance written in their native languages without presupposing that this utterance is read, which requires them to be involved in perceptually decoding the alphabetic writing, either Latin or Cyrillic, and mapping those into complete words, then these into complete sentences etc., which is itself a very complicated process that can be disrupted in a number of ways (Dehaene, 2009).

All the steps of the process of reading rely on processing physical vehicles of information in particular locations of the brain, including areas responsible for so-called visual word form area (VWFA, which need not be at all strictly visual; cf. (Siuda-Krzywicka et al., 2016)). To summarize, the mere fact that we do not usually have to mention these cognitive operations in order to understand a simple story about a Frenchman killing his uncle does not necessarily mean that they do not happen. The intuition pump is simply misleading in this regard. For all we know, these operations must occur for the story to hold true, and they require information to be physically encoded. If it had not been physically encoded, it would not exist.

And if this information is encoded, then it could be measured in multiple ways. The simplest one is just to count the number of symbols in the encoding (the logical depth of the message, or the number of degrees of freedom).

Let me, then, summarize what Dennett's intuition pump could establish. It points out that one could have differing pieces of semantic information, whose satisfaction conditions could be, to some extent, the same. Namely, Jacques, Sherlock, Tom, and Boris entertain thoughts which would be false if no Frenchman killed his uncle in Trafalgar Square. But the intuition pump does not require us to think that semantic information could be transmitted or saved without being encoded.

4. Encoding and measuring semantic information

Let me put Dennett's account in a larger context. At first glance, the idea that one could have truth without encoding it seems to violate a very basic assumption of all semiotics, namely that signs have vehicles. To say that one could have a true propositional sign that does not have any vehicle is to commit a straightforward contradiction. Signs have vehicles by definition. Moreover, for any physicalist, vehicles are physical.

While codes and encoding are criticized as wrong-headed ways of thinking about cognition and the brain (Bickhard, 1993; Brette, 2019), these criticisms stem from a certain misapprehension of what encoding is supposed to do and how it relates to an old idea of the correspondence theory of truth. The major criticism of applying the notion of encoding to cognitive representation is that encoding, as they say, cannot supply the meaning of representation, which must come from elsewhere. Critics seem to imply that meaning has to be established first, before it is ever encoded. This is not exactly true, and such projects as distributional semantics in linguistics or s-vector semantics (Isaac, 2017; Shea, Godfrey-Smith, & Cao, 2018; Skyrms, 2010) attest that mere connections between vehicles of information could be meaningful, even if they cannot fully determine the referential aspect of meaning. But most basically, a simple mapping, whether in the case of simple indicators or structured representations, between a physical vehicle of information and its reference, can be cognitively used by

us only if we also know how to find structure in the referential domain. In other words, an idea that a proposition encodes reality requires that we are able to treat reality at least somewhat like an information vehicle.

This may seem like an absurd idea. But there have already been ambitious projects to understand carving reality at informational joints in a fully naturalized way, without presupposing a God's-eye-view on reality, with neatly delineated and labeled kinds and individuals. Nonetheless, it is not a coincidence that Tarski denied that his approach to truth could be easily extended to cover natural language (White & Tarski, 1987); not only natural languages give rise to unavoidable paradoxes but they also are extremely difficult to analyze in terms of what satisfies their constituent expressions, which, in contrast, can be done axiomatically for a formal system. 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.

The core assumption of Dretske's account of semantic information (Dretske, 1982) and of situation theoretic framework (Israel & Perry, 1990) is that objective information is currently available. This makes it easy for them to talk of how information could be true of structured reality. Satisfaction conditions are then simple to define in objective terms. This is, of course, difficult for pragmatists and constructivists to swallow, as the whole project of finding physical information smacks of objective meaning as existing "out there". But naturalized semantic information does not require that we ignore biological agency. If we look at these projects not as *opposed* to but *complementary* to attempts to see biological agency in the world, then it is easy to see that any attempt to understand it has to, sooner or later, refer to control theory. And there is a well-known result in control theory, the good regulator theorem (Conant & Ashby, 1970), which states that the optimal regulator contains an isomorphic model of what it controls. This, surely, presupposes that the plant (as control theory calls the controlled entity) has some structure, to which the model may be isomorphic. (Note that there could be non-optimal regulators, whose models are not as detailed or more detailed). In other words, we have to find ways to find causally relevant structure in controllers and plants in order to use control theory. If we do not assume this much, we cannot really use control theory. And without control theory (or something equivalent to it), what can we state about agency? Not much. And what we will state will not even be remotely as precise.

It suffices to ascertain a certain structure in a physical vehicle to establish that (structural) information is encoded by it. Whether the structure is continuous or not, we can always measure the structural information content (the number of logons) in the physical media. For control systems, the only structural information that counts is that which is causally relevant for the functioning of such systems. And having defined encoding as causally involved in the structure of a control system, we can then apply control theory to our system.

MacKay was onto something: there is a connection between information and control, and this connection was brought to the fore by Conant and Ashby in their somewhat obscure and diffi-

cult treatment of this problem.¹ Rejecting the idea of encoding is then simply premature. The costs of such a rejection are rather prohibitive.

But as soon as we get some encoding, we may begin to think of measuring information. Still, the problem of measuring semantic information remains difficult. One could follow Gabor and measure information in terms of its logical depth, but it does not track the semantic aspect of information, just its structure. One could also try to measure the logical depth of the information's satisfaction conditions. Alternatively, could follow Carnap and Bar-Hillel, and measure the information yield of a statement within a given language in terms of the set of possible states it rules out and a logical probability space over those states (D'Alfonso, 2011). Some even argue that a fair amount of our intuition about informativeness of a message is already present in Shannon's entropy (Isaac, 2017), sometimes connected with the idea of autonomous dynamical systems (Kolchinsky & Wolpert, 2018).

These are serious projects and should not be too easily dismissed. While, today, information theorists are (fortunately!) not so eager, as they were in the 1960s or 1970s, to estimate the informational capacity of the brain, at the end of the day we have to get around to quantitative approaches to mental functioning. This is also why measuring consciousness is one of the most important successes in the recent history of the scientific study of consciousness (Boly et al., 2013; Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008). Without it, all experimental work cannot simply disconfirm strict theoretical predictions.

5. Conclusion

Dennett's idea of semantic information as design worth getting is surely interesting, and it would be fruitful to see how similar it is to new accounts of functional information (Fresco, Ginsburg, & Jablonka, 2018). But his argument fails to establish that design worth getting floats free from any encoding. And as soon one has to talk of encoding, one can talk of measurement.

In this commentary, my aim was to provide an alternative analysis of the intuition pump used by Dennett. All it does is point out that information could be encoded in various ways and retain the same satisfaction conditions. This is hardly news, and cannot substantiate Dennett's strong claims of independence of semantic information of encoding. Semantic information is always encoded and thereby its measures are definable.

¹ One thing that could escape one's attention, when reading their paper, is for example that the plant need not contain anything external to a control system (no external feedback loop is required), thus control may be limited to the endogenous dynamics of a physical system. The good regulator theorem does not establish the necessity for representations of external reality for any controller. It only establishes it for systems whose plant is defined in such a way as to encompass that part of reality.

Acknowledgements

I wish to thank Paweł Gładziejewski, Tomasz Korbak, and Michał Piekarski for their helpful comments on the draft of this contribution. The work on this commentary was funded by a National Science Centre (Poland) research grant under the decision DEC-2014/14/E/HS1/00803.

References

- Bickhard, M. H. (1993). Representational content in humans and machines. *Journal of Experimental & Theoretical Artificial Intelligence*, 5(4), 285–333. doi: 10.1080/09528139308953775
- Boly, M., Seth, A. K., Wilke, M., Ingmundson, P., Baars, B. J., Laureys, S., ... Tsuchiya, N. (2013). Consciousness in humans and non-human animals: Recent advances and future directions. *Frontiers in Psychology*, 4. doi: 10.3389/fpsyg.2013.00625
- Brette, R. (2019). Is coding a relevant metaphor for the brain? *Behavioral and Brain Sciences*, 42, E215. doi: 10.1017/S0140525X19000049
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97. doi: 10.1080/00207727008920220
- D'Alfonso, S. (2011). On Quantifying Semantic Information. *Information*, 2(1), 61–101. doi: 10.3390/info2010061
- Dehaene, S. (2009). *Reading in the Brain: The New Science of How We Read*. Penguin.
- Dennett, D. C. (1969). *Content and Consciousness*. London: Routledge and Kegan Paul.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Mass.: MIT Press.
- Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.
- Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. London: Allen Lane.
- Dretske, F. I. (1982). *Knowledge and the Flow of Information* (2nd ed.). Cambridge, Mass.: MIT Press.
- Floridi, L. (2010). *Information: A very short introduction*. Oxford: Oxford University Press.
- Fresco, N., Ginsburg, S., & Jablonka, E. (2018). Functional Information: A Graded Taxonomy of Difference Makers. *Review of Philosophy and Psychology*. doi: 10.1007/s13164-018-0410-7
- Gabor, D. (1946). Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93(26), 429–441. doi: 10.1049/ji-3-2.1946.0074
- Harms, W. F. (1998). The Use of Information Theory in Epistemology. *Philosophy of Science*, 65(3), 472–501. doi: 10.1086/392657
- Isaac, A. M. C. (2017). The Semantics Latent in Shannon Information. *The British Journal for the Philosophy of Science*. doi: 10.1093/bjps/axx029

- Israel, D., & Perry, J. (1990). What is Information? In P. Hanson (Ed.), *Information, Language and Cognition* (Vol. 1, pp. 1–19). Vancouver: University of British Columbia Press.
- Kolchinsky, A., & Wolpert, D. H. (2018). Semantic information, autonomous agency and non-equilibrium statistical physics. *Interface Focus*, 8(6), 20180041. doi: 10.1098/rsfs.2018.0041
- MacKay, D. M. (1969). *Information, mechanism and meaning*. Cambridge: M.I.T. Press.
- Miłkowski, M. (2015). The Hard Problem Of Content: Solved (Long Ago). *Studies in Logic, Grammar and Rhetoric*, 41(1), 73–88. doi: 10.1515/slgr-2015-0021
- Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: Relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, 12(8), 314–321. doi: 10.1016/j.tics.2008.04.008
- Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(July, October), 379–423, 623–656.
- Shea, N., Godfrey-Smith, P., & Cao, R. (2018). Content in Simple Signalling Systems. *The British Journal for the Philosophy of Science*, 69(4), 1009–1035. doi: 10.1093/bjps/axw036
- Siuda-Krzywicka, K., Bola, Ł., Paplińska, M., Sumera, E., Jednoróg, K., Marchewka, A., ... Szwed, M. (2016). Massive cortical reorganization in sighted braille readers. *ELife*, 5(MARCH2016), 1–26. doi: 10.7554/eLife.10762
- Skyrms, B. (2010). *Signals: Evolution, learning, & information*. Oxford, New York: Oxford University Press.
- White, M., & Tarski, A. (1987). A philosophical letter of Alfred Tarski. *The Journal of Philosophy*, 84(1), 28–32.

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.