Review of: "[Paper removed, will be broken down to two different projects] Big Data, Granger Causality Analysis, and the Undecidability Property of Neuroimaging"

Farid Zahnoun

Potential competing interests: The author(s) declared that no potential competing interests exist.

Review of P.F. Bloniasz: 'Big Data, Granger causality analysis, and the undecidability property of neuroimaging' (2022)

Before starting my commentary, I would first like to clarify that I am a philosopher, not a neuroscientist. As such, I only feel sufficiently qualified to evaluate the conceptual, logical and argumentative issues presented in this paper. I will leave it to my scientifically trained colleagues to comment on the more empirical or mathematical sections of this work. This being said, I think it's safe to say that the bulk of this paper precisely deals with conceptual issues.

Unlike most theorists working in neuroscience, and much to his credit, the author decides to take a more philosophical-critical approach. Such an approach is, in my opinion, much needed. In fact, if I had to voice my worries with this paper in one sentence, it's that the paper is not critical enough, especially in its analysis of the notion, or rather: *notions* of information that are being discussed. The author is certainly right, both in his critique of the so-called 'end of theory view' and in his emphasis on the principled undecidability of neuroimaging. This, in and of itself, makes the paper a valuable contribution to the field. At the same time, however, the author's criticism stays well within the boundaries of the informational paradigm, where the idea that the brain is essentially an information processing device is accepted as a given. This also remains the central background assumption of the present paper, notwithstanding the fact that the author realizes that the notion of information "requires a bit more nuance than neuroscience has given it historically". Nevertheless, the paper reiterates — and thereby reinforces — some of the same flaws that characterize almost every neuroscience paper that invokes the notion of information: the conflation between intentional or semantic information (information as meaningful messages) with Shannon's statistical and quantitative notion of information; second, the objectification of information (the role of the receiver is non-constitutive) ; third, the confusion between the perspective of the experimenter and the perspective of the brain itself (if sense can be made of this idea). To keep this commentary at a reasonable length, I will limit myself to a discussion of the first issue, which I think is the most important.

Intentional information and Shannon information

Ever since Shannon's 1948 publication of "A Mathematical Theory of Communication", theorists, including Shannon himself, have been emphasizing that his notion of information has in itself nothing to do with what we ordinarily would call information, i.e., some meaningful and truth evaluable proposition (e.g., a message) which increases our knowledge. Below are just a few earlier quotes[1] from scientists and philosophers who have emphasized the crucial distinction between meaningful information on the one hand, and Shannon's technical notion on the other.

In discussing Information Theory, Fritz Machlup writes in his co-edited 1983 book *The Study of Information: Interdisciplinary Messages*:

[T]o speak of information when information theorists explain their system is a sad misuse of language; they explicitly abstract from a meaning-content of the signals the transmission of which they describe. Their system does not care about telling anything, directing or advising anybody, arousing anybody's interest, or inducing any decisions or emotions. Appropriate words to use in the context would be signal transmission, actuation, or activating impulses. The use of the word information in this sense has led to unending confusion and should no longer be condoned. (Machlup 1983: 660)

And two pages earlier, Machlup writes: "They [*the information theorists, m.a.*] use information and amount of information in a sense that has so little to do with any traditional or metaphoric meanings of the word that one can only wonder why the scientific community has allowed it to continue." (Machlup 1983: 658)

Indeed, using the same word for two very different things is bound to cause confusion, something which was also realized by Bar-Hillel, who, together with Carnap, developed a theory of semantic information. Their theory is explicitly contrasted with Shannon information. Bar-Hillel argues that Shannon's information theory had better be named 'Theory of Signal Transmission'. He adds:

Even more important than the change of name from Information Theory to Theory of Signal Transmission (plus, perhaps, Theory of Coding) would be to discard the use of the term information' within this theory, with all its ambiguities and semantic traps. (Bar-Hillel 1973: 134, fn. 22)

That Shannon's theory was bound to cause confusion was already realized by Joseph C. R. Licklider. As early as 1951, he writes: "It is probably dangerous to use this theory of information in fields for which it was not designed, but I think the danger will not keep people from using it" (quoted in Kline, 2015: 58). What, precisely, is this danger Licklider and others are warning us about? The danger is that we fall prey to the following fallacy: because we can apply Shannon's information theory to a certain phenomenon (neural activity, say), that which we are applying it to must necessarily involve the sending, receiving and coding of informational messages, just as Shannon applied his theory to the sending, receiving and coding of informational messages. This, of course, is wrongheaded. The mere applicability of Shannon's theory proves nor disproves the presence of meaning, nor that of a message, for that matter. Shannon's communication theory *assumed* or *presupposed* information-carrying messages. But this does not mean that, just because we can apply his theory to other domains, this warrants the assumption that there is any sending, receiving or coding of messages happening here as well (one can use Shannon's theory to measure how much information the role of a pair of dice carries, but that hardly warrants the assumption that the dice are sending, receiving or coding for a message). In other words, the fact that we can use Shannon's theory within the domain of neuroscience does not in any way warrant the very different supposition that neurons are in the business of sending, receiving and encoding meaningful messages. In fact, as has been argued, we have good reasons for not accepting this assumption. One good reason is that both ordinary intentional information and Shannon information constitutively involve a receiver/observer/interpreter/informee which must already possess specific cognitive abilities which can't be attributed to neurons or brain regions, at the risk of committing the mereological or homuncular fallacy (see, for instance, Bennett & Hacker 2003). In case of Shannon information, for instance, there is a constitutive dependency on the epistemic states of the receiver: he or she must *know* the probability distribution of the possible outcomes. How does this apply to neurons or brain regions? Yet, without such relation to the epistemic states of the receiver/observer, there simply is no Shannon information (which is why it *constitutively* involves a receiver/observer). In any case, that Shannon's notion of information is not to be confused with ordinary semantic information was made clear by Shannon himself already in the opening lines of his original 1948 paper, he writes:

The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. *These semantic aspects of communication are irrelevant to the engineering problem*. The significant aspect is that the actual message is one selected from a set of possible messages. (Shannon 1948: 379, m.i.)

And three years later, E. Colin Cherry writes: "It is important to emphasize, at the start, that we are not concerned with the meaning or the truth of messages; semantics lies outside the scope of mathematical information theory." (Cherry 1951: 383). To this, Carnap & Hillel add:

It has, however, often been noticed that this asceticism is not always adhered to in practice and that sometimes semantically important conclusions are drawn from officially semantics-free assumptions. In addition, it seems that at least some of the proponents of communication theory have tried to establish (or to reestablish) the semantic connections which have been deliberately dissevered by others. (Carnap & Hillel 1952: 2)

What applied to "some of the proponents of communication theory" in 1952 applies, it seems safe to say, to the vast majority of present day neuroscience theorists, including the author of the present paper. The problem is always the same: the fact that a cognitive observer can correctly say that, from his or her perspective, there is Shannon information in the brain does not mean that there literally is interpretable semantic or, as I prefer to call it, intentional information in the brain. I use 'intentional' here in the classic phenomenological sense of 'aboutness'; information, then, that is about something. The idea that the brain is processing this latter kind of information — which is widely accepted by neuroscientists and cognitive scientists alike — is a wholly separate idea which is in no way made more plausible by the applicability of Shannon's statistical theory of quantifying, what he decided to call, information. Indeed, as we've just seen, many authors agree that this terminological choice is a historical blunder.

How does all of this apply to the present paper? On the one hand, at times, the author correctly describes Shannon information in terms of transfer entropy and, at least at the outset, makes sure to safeguard Shannon information from semantic or intentional connotations. In fact, the author's decision of sometimes using scare quotes around 'information' seems to indicate that the term 'information' is to be used with caution here. However, at the beginning of section 2, the author drops his guard. He starts out with presenting us with a description of information, a description which is "often implicitly or explicitly stated" to be commonsensical: "information is the set of properties of a signal waveform produced by a neurobiological process that can be measured and characterized, for example, by its covariance to other signals." The author continues: "Regardless of the context, the 'information' in the signal is, supposedly, operationally clear: we are studying action potentials, the summative electrical dipoles of extracellular activity, or some sort of ionic behavior that underlies the action potential (e.g., Na+/K+)." It is entirely unclear how studying action potentials or other neural behavior relates to the idea of there being information carrying signals. I assume that all the author is saying here is that the information just is what the experimenter learns through her measurements of neural behavior. This, of course, is an entirely trivial use of the notion of information, and it has absolutely nothing to do with the completely different idea of the brain as an information processing organ. The author goes on: "Some more concerned with the details of information processing in neuroscience might problematize the obviousness of 'information' by pointing out the various ways it might be encoded in the brain." What is allegedly encoded by the brain, however, is not Shannon information, but some (neurobiological) message which carries a kind of informational content. The only way to make any interesting sense of the idea that the brain is coding for information is to assume that this information is intentional. It has to be *about* something (for instance, changes in light intensity). It is also in this sense that neuroscientists use 'information' in connection to coding. (See Brette 2019) The author then follows de-Wit et al. in presenting 6 ways in which information might be carried in the brain. But again, the only way to make any interesting sense of this idea is not in terms of Shannon entropy, but in terms of intentional messages: it is based on the content of these coded messages that the brain is thought to operate. And this must be the kind of information the author has in mind when he writes that the fact of there being many ways in which information might be carried is "not too daunting as, most plausibly, there is some sort of relevant information across each of these." The author then gives an example:

When looking at the cardiac ganglion of the *Homarus americanus*, we might be interested in any number of features in a neurophysiological waveform from extracellular cell recordings: the duty cycle, spiking frequency, driver potential and/or spike amplitude, and so forth help with understanding both the underlying dynamics of ionic channels and the emergent behavior of the network. (e.g., Cooke, 2002). The challenge comes from isolating each and seeing which waveform feature 'says' what within a given research paradigm, or whether a given characteristic appears consistently but is merely noise. Determining how information is being propagated is difficult, and is not purely a computational problem…

Two very different ideas are being lumped together here. On the one hand, there is the correct but trivial idea that information is that which "helps with understanding". But, again, this has nothing to do with information in the brain, it is simply information *about* the brain. On the other hand, however, it is clear enough that this very same information about the brain is assumed to also exist at the level of the brain itself, where it is said to be "encoded" and "propagated". But why would the information we gain from studying duty cycles, spiking frequencies, driver potentials or spike amplitudes say anything about the alleged information being processed by the brain, or whether there is any processing, encoding or decoding of information going on at all?

However, it is especially in his discussion of Information Theory (in section 2.2) that the conflation between intentional information (messages) and Shannon information becomes apparent. The author starts with Shannon's own definition of information: "information is the reduction in uncertainty obtained when a selected outcome of interest, x, from the probability distribution, p(x), is observed (Shannon, 1948; Wibral et al., 2014). This is called Shannon information." Note again that Shannon information by definition constitutively involves an observer. A few lines further, we read the following rather puzzling lines: "if an outcome, x, is relatively improbable, more information is transferred given the event is actually observed and observed consistently. This is an intuitive result, because a non-probable event occurring and occurring consistently can be thought of anthropomorphically as some biological system intentionally sending information to some other biological system." The author correctly emphasizes the importance of the observer. It is unclear to me, however, how an event can both be non-probable, while at the same time occurring consistently. I would think that there is precisely less Shannon information being transferred, not more. Also, I fail to see how thinking in terms of a biological system equates to an anthropomorphism, although the following example of the noisy dinner party is certainly anthropomorphic. It is also in this example that we find another clear conflation between Shannon information and meaningful, semantic information:

If person 1 shouts out the name of person 2 in person 2's direction multiple times and person 2 hears this call, this is a relatively low probability event that is being observed multiple times (i.e., *it is more meaningful* and, thus, more information is being spread to person 2 relative to the background noise in the room).

The idea that 'more meaning' implies more Shannon information is mistaken and violates one essential property of Shannon information highlighted by Shannon himself, and, as we have seen, countless others, i.e., that it abstracts away from meaning (cf. supra). And the same misinterpretation occurs several other times in what follows, for instance in these lines: "Shannon's information might be a message that is clear to the experimenter, but there is no guarantee, or even a high probability guess, that the potential information that is selected by the experimenter is actually a meaningful signal." Again, Shannon's information is not, and has nothing to do with the meaning or the meaningfulness of a message, nor is it in itself identifiable with a message. So it also has nothing to do with interpretation. Yet, on several occasions, the author nevertheless claims that Shannon information is a kind of information that can be interpreted. Again following de-Wit et al., the author writes: "Barrett et al. (2018a, 2018b) are making a classic mistake in interpreting Shannon's communication system; realistically, "there is no such thing as objective information, because it is always subject to interpretation by a receiver" (de-Wit et al., 2016, p.1417)." The quote by de-Wit et al. is in itself philosophically interesting, but that discussion will have to wait for another day. What I am interested here is the connection with Shannon information and 'being interpretable'. Again, ordinary semantic or intentional information is indeed interpretable, i.e., the notion of interpretation logically pertains to this kind of information. It does not, however, logically pertain to Shannon information. Shannon entropy simply isn't the kind of thing to which interpretations apply, because Shannon entropy is a purely quantitative notion and has nothing to do with interpretable meaning. Interpretation, on the other hand, has everything to do with meaning. Moreover, the author not only conflates Shannon information with interpretable semantic information, he presents it as a sufficient condition for Shannon information:

Target information (i.e., what is encrypted) will appear as noise to any receiver that does not have the correct decryption key. If someone has the correct decryption key, *then the information is interpretable–therefore making it information by Shannon's definition*.

Claiming that something has to be interpretable to count as Shannon's definition goes radically against Shannon's own ideas. At no point did Shannon define information in terms of interpretation or interpretability. Quite the contrary. And it certainly doesn't figure in the definition the author has presented earlier (i.e., information as the reduction of uncertainty).

Returning to the idea that the selected information by the experimenter may or may not be a "meaningful signal", from the perspective of Information Theory, it is entirely irrelevant whether the experimenter is observing a *meaningful* signal or not. It is, however, very relevant from the perspective of the neuroscientist who assumes that there is *meaningful*, non-Shannon information in the brain. Clearly, the kind of information the neuroscientist speaks of when she is talking in terms of the brain sending, receiving, and coding for information is precisely *not* Shannon information, but information which is not just akin, but identical to the content of the meaningful messages we (i.e., persons) send, receive and code for. And it is also this notion of a meaningful message which takes center stage in the author's so-called 'undecidability property of neuroimaging, or UPN for short. As the author puts it:

UPN claims the following: given a neurobiological message that is described by stochastic time process X, (t), and is produced from an information source, I, in closed neurological system, S, if the message from I in S is completely measured in a channel, then it is necessarily true that the following decision problem is undecidable. In S, is the received signal, Ω, which is decoded from X(t)–given the signal is known to flow to some receiver, ϕ , in destination, D–the exclusive information sent from I to drive some function, F?

The central notion here is that of a (neurobiological) message, which is said to be measurable. Again, then, the author falls back on the intentional notion of information. There are said to be messages within a closed neurological system (so not for the observing experimenter), which must be *about* something (a message that is not about something is not a message). If this should somehow be connected to Shannon's idea of 'measuring information', I don't see how. Shannon's theory is not at all about measuring messages, if there could ever be such a thing (how does one measure a message?). What the author seems to be saying is simply that, *given the assumption that the brain is sending, receiving and coding for meaningful messages*, as a matter of principle, it is impossible of deciding what the content of the messages is based on measurements provided by neuroimaging techniques. The author, however, puts it differently:

In a less formal sense, UPN merely claims that neuroimaging, intrinsically, cannot exhaust all criteria required to determine the information flow from one brain region to another–more empirical work, consistent with theory to determine the code, needs to be done outside of the neuroimaging approach. The reason is because, as mentioned, neuroimaging does not directly measure information.

Again, from a Shannonean information theoretic perspective, it is unclear what directly or indirectly measuring information (in the sense of the information of a message) could possibly mean. What is measured is the *amount* of information, understood in terms of uncertainty reduction, relative to a known probability distribution. And although the author never explicitly speaks in terms of the informational content of the neurobiological message, it seems clear enough that this is the kind of information (not Shannon information) he has in mind. What else is there to code for if not informational content? And what else is there for the other brain regions to "decode to complete the specific brain task" or "to drive some function, F"? As said, all of this already assumes that the brain really is in the business of sending, receiving and coding for (neurobiological) messages. According to the author, this is a known fact: "researchers using GGC know very well that the brain regions they are looking at are communicating and are, in fact, communicating in a particular code." I know of at least one excellent paper by an outstanding neuroscientist who challenges this 'fact' (See Brette 2019: *Is coding a relevant metaphor for the brain?)* Perhaps brain regions are communicating in *some* sense (in the sense of communicating vessels, perhaps. See Cherry 1957: 219 for further discussion of proper vs. improper uses of the term 'communication'), but whether they are communicating in the sense of sending messages which carry intentional information is far from settled. More importantly, what is also still far from settled is how, assuming that there are indeed meaningful messages being transferred in the brain, this meaningful information is "used" to "complete a specific brain task", or how it manages to "drive" a brain function? How does information come to play a causal role within a system? We know that the meaning of a message makes a difference to us, human beings. If somebody tells me that my mother just died, this will have a tremendous causal effect on me. But how does this work at the level of neurons? Neuroscience often proceeds as if this issue has already been solved. To my knowledge, it hasn't. In fact, it is one of the biggest outstanding issues in philosophy of mind: how can the content of my mental states have causal effects? Apparently, what counts as a challenging problem in one discipline is taken to be so unproblematic in another field that it can be invoked as an *explanans* instead of an *explanandum*.

As said at the beginning, I think the paper raises a few important issues within the boundaries of the dominant informational paradigm. At the same time, however, it inadvertently covers up bigger worries with some core assumptions of the informational paradigm itself by paying insufficient attention to a few fundamental conceptual distinctions. In the above, I have tried to clarify the importance of one such distinction. If theorists keep confusing Shannon's statistical notion of information with the more commonsensical, but also more homuncular notion of intentional information, the very idea of the brain-as-information-processor is bound to remain equally confused.

References

- Bar-Hillel, Y. (1973). *Language and Information.* London: Addison-Wesley.
- Bennett, M.R. & Hacker, P.M.S. (2003). *Philosophical Foundations of Neuroscience*. Wiley-Blackwell
- Brette, R. (2019). Is coding a relevant metaphor for the brain*? Behavioral and Brain Sciences, 42*: 1-44.
- Carnap, R. & Bar-Hillel, Y. (1954). An outline of a theory of semantic information.*Journal of Symbolic Logic 19*(3):230-232.
- Cherry, E.C. (1951). A history of the theory of information.*Proceedings of the Institution of Electrical Engineers 98* 383 – 393.
- Cherry, E. C. (1957). *On Human Communication*. MIT Press.
- Deacon, T. (2011). *Incomplete Nature*. W. W. Norton Company
- Godfrey-Smith, P. (2012). Review of Brian Skyrms' *Signals*. *Mind 120*: 1288-1297.
- Kline, R. R. (2015). The cybernetics moment. Baltimore, MD: Johns Hopkins University Press
- Machlup, F. (1983). Semantic quirks in studies of information. In F. Machlup & U. Mansfield (Eds.), The study of information: Interdisciplinary messages (pp. 641-671). New York: Wiley.
- Nizami, L. (2019). Information theory is abused in neuroscience.*Cybernetics and Human Knowing. Vol. 26* (4): 47-97.
- Searle, J. R. (2013). Can information theory explain consciousness? Review of*Consciousness: Confessions of a Romantic Reductionist* by Christof Koch. New York Review of Books.

Shannon, C. (1948). A mathematical theory of communication.*Bell System Technical Journal 27*: 379-423.

- Shannon, C., & Weaver, W. (1949/1972). *The Mathematical Theory of Communication*.

Urbana, IL: University of Illinois Press. (Original work published in 1949)

[1] For more recent references to authors emphasizing the essential distinction between Shannon information and meaningful information, see, for instance, Godfrey-Smith 2012, Deacon 2011, Searle 2013, Brette 2019. See especially Nizami 2019 for an extensive treatment of how Information Theory is being “abused” in neuroscience.