Information and representational revisionism

Whit Schonbein

Department of Computer Science
University of New Mexico, Albuquerque

whit.schonbein@gmail.com

DRAFT – DO NOT QUOTE
2014-04-17

ABSTRACT: While information theory (IT) is generally viewed as philosophically unproblematic, theories of mental representation built upon it are not: Representational revisionist philosophers increasingly call on cognitive science to divest itself of such 'indicator' notions of representation. Given the importance of IT to cognitive science, it is reasonable to inquire into its status under these critiques. I argue that current revisionist arguments have the bases covered: (1) Ramsey's (2007) argument implies that the use of IT should be curtailed on the grounds that at least some devices presumed to process information in fact do not have that function; (2) Hutto & Myin's (2013) approach makes room for IT *sans* content; and (3) the fate of IT under Chemero's (2009) approach depends on how details about the epistemological role of IT are spelled out. These results also show that revisionist arguments, even when limited to the context of information theory, need not be 'merely verbal', as some have charged. Instead, they can have radical consequences for how we understand cognitive processes, i.e., not only sans representation, but also *sans information*.

1. Introduction

According to 'indicator' theories of representation, a state of a system qualifies as a representation just in case it (i) carries information (by virtue of causally co-varying with states of other components or the environment), and (ii) has the function of carrying that information, i.e., the component was recruited into the system because of the fact it 'indicates' certain conditions obtain. For instance, the state of the bi-metallic strip in a mechanical thermostat represents the temperature of the ambient atmosphere because its state causally co-varies with ambient temperature, and it is used by downstream components for its capacity to carry that information (Dretske 1981, 1988; Millikan 1989; Fodor 1990).

Unfortunately for indicator representation, a growing chorus of voices call for the divestment of the notion from the cognitive science toolkit.[1] For instance, Hutto & Myin (2013) claim that "[c]ognitive systems don't "pick up" or "take in" any informational contents; there are no such things as informational contents to take in," (p. xvi). Likewise, Ramsey (2007) claims that "the [indicator] notion is too weak to have any real explanatory value," (p. xvii) and that "it is not a useful theoretical posit of cognitive science" (p. 118). Finally, in a similar vein, Chemero (2009) argues that "the representational story could be told but [it's] not particularly relevant" (p. 72). Despite their differences, these authors share a common commitment to limiting or eliminating appeals to indicator

1 The term 'anti-representationalist' is sometimes used to describe this sort of thesis; here I use 'representational revisionist' because not all of the authors discussed here go so far as to endorse the elimination of all representation-talk (e.g., Hutto & Myin 2013, p. xviii).

representation in the cognitive sciences.

In this paper I consider the question: *What happens to information theory* (IT) *under these arguments?* This question is important for at least two reasons. First, IT is indispensable to an epistemically thorough cognitive science.[2] Consequently, given the conceptual proximity of indicator representation to IT, when researchers attack the former it is important to also assess the potential impact on the latter. Second, an inevitable response to revisionist critiques is some version of the 'merely verbal' objection; Ramsey (2007) summarizes it as follows:

> Perhaps there are no deep differences between [non-representational] things like firing pins or immune responses on the one hand, and states we call receptor representations on the other. But what is the harm in referring to the latter as representational states? If doing so is a conceptual confusion, then it appears to be a conceptual confusion that has no negative consequences. (p. 146)

With respect to IT, the issue is this: What does calling a state 'representational' have to do with the construction of information-theoretic models? Information theory is widely perceived as "unproblematic" (Godfrey-Smith 2007) or even "impoverished" (Floridi 2004) precisely because it does not deal with content. So it is easy to assume that critiques of indicator representation leave IT intact. While researchers developing IT models certainly rely on there being nomic relations between components, and perhaps also presuppose that these causal relations are related to the functions those components play in the target system, it is not clear that the representation- or content-talk that often accompanies these practices involves any substantive appeal to semantic features of the sort representational revisionists find so objectionable. Consequently, if the goal of revisionist critiques is to suppress representation-talk, then doing so does not seem to have any consequences for that part of the account that is doing the 'real' epistemological work, namely, information theory. In contrast, if revisionist arguments also impact the use of IT, then it is hard to view those arguments as 'merely verbal'.

With this in mind, after offering some basic background information on information theory, I argue that Ramsey's (2007) critique has significant consequences for how information theory is deployed in cognitive science (it is not 'merely verbal'), Hutto & Myin's (2013) does not (it is 'merely verbal'), and the status of information theory under Chemero's (2009) approach depends on how his critique is extended – a project I sketch but whose completion is beyond the scope of this paper.

To be clear, my goal in this paper is not to defend indicator representation, but rather to assess the status of information theory under those critiques. This occasionally requires adopting a critical perspective, but in doing so I am not trying to rescue indicator representation. Similarly, in the event a revisionist argument implies a radical rethinking of how we use information theory, it may be tempting to interpret such a result as a *reductio* of that argument; again, I try to remain neutral in this regard. The overall goal is to assess whether, from the perspective of information theoretic modeling, representational revisionism can be dismissed as a 'merely verbal' dispute.

2. Information Theory...

In this section I use three examples to (i) illustrate the epistemological utility of information theory (IT), and (ii) identify some broad presuppositions of that theory (so that we can appreciate how it might

---

2 See, for instance, Milenkovic (2010) and Dimitrov (2011), and the articles contained in those volumes.

be challenged by revisionist argument).[3]

The first example is a simplified gun: When the trigger is depressed, a firing pin slams against the back of a shell casing, causing a chemical reaction that forces a projectile through the barrel and into the air. We can concoct a mathematical description of this causal chain, from force applied to the trigger, transferred via some mechanism to the firing pin and to the shell casing, resulting in the deformation of that casing, which impacts the gunpowder, etc. Call this hypothetical causal story the 'C-description', 'C-explanation', or 'C-story' of the gun's behavior.

The tools of information theory provide a complementary perspective: Whereas the C-story describes the sequence of events and interactions that occur given the state of the trigger, the information-theoretic account quantifies the probabilistic relations holding between possible states of the components involved. For instance, we know from the C-story that the trigger can take on two possible states, depressed (T_A) or not (T_B), and at any given moment it actualizes one of them. The quantity of information carried by these states is a function of their probability: The less probable an event, the more information it carries. Metaphorically speaking, one can think of information as a measure of the 'distance' a state must travel from possibility to actuality; less likely states travel greater distances, and hence their actualization carries greater information. By convention we use logarithms to characterize this distance. For instance, if each state is equally probable, they travel the same distance – ½ of the sum of the two distances – so both T_A and T_B carry $-\log_2(½) = 1$ bit of information, and the set of states instantiated by the trigger carry 1 bit on average. In contrast, if T_A occurs ¼ of the time while T_B occurs ¾ of the time, T_A carries $-\log_2(¼) = 2$ bits of information while T_B carries $-\log_2(¾) = 0.415$ bits of information, with the overall set carrying 0.8112 bits on average.[4]

With an information metric for the set of possible states of a single component in hand, we can ask how those of different components are related, i.e., how much information is *shared* across states of multiple components. For instance, the firing pin of our gun can also be in two states, down (P_A) or up (P_B). Assuming they are equiprobable, each state carries one bit of information. We know from the C-story that there are nomic dependencies holding between the state of the pin and the state of the trigger, so we can ask: To what degree does the fact the pin is in state P_A or P_B imply anything about the state of the trigger? In other words, how much of the one bit of information carried by a state of the pin is determined by the state of the trigger, and how much is due to noise? That is, how *reliable* is the relation between the state of the trigger and the state of the trigger?

There are three possibilities: If there is no nomic connection between the two, then the state of the pin is unrelated to the state of trigger outside of a happy accident; in this case the *mutual information* shared between trigger and pin is zero. If there is such a connection, but it is corrupted by noise, then some part of the one bit of information carried by the state of the pin is 'inherited' from the state of the trigger, with the remainder going to outside influences. Finally, if the relation is perfect, then the entire bit of information carried by the pin is the result of the state of the trigger; the mutual information shared between pin and trigger is optimal, and one could infer without error from the state of the pin to the state of the trigger.

This is, in a nutshell, what IT does: It takes the causal processes of a device and, beginning with the states of single components, quantifies the relations between changes in states across the components of those devices given the probabilities of those states being actualized. In this way, information theory can tell us (among other things) how much influence a change in state of one component has on a change in state of another components, and hence how reliable that relation is.

3 The survey of IT offered here is cursory, but we don't need much for present purposes; for a respectable overview see XXXXXXXX.

4 This description deliberately emphasizes the objective (rather than subjective) approach to information; see Thornton (2013).

That being said, one issue with the first example is that guns are typically interpreted as 'merely causal' mechanisms: Firing pins and triggers may carry information, but that is not their “job” (Ramsey 2007), so applying information theory to that case is misplaced. So let us introduce a second example where it makes more intuitive sense to apply IT: The mechanical thermostat. In such a device there is a bi-metallic strip coiled into a spiral and as the ambient energy in the room changes the metals forming the strip expand (or contract) at different rates, causing the spiral to ‘rotate’. A mercury switch is attached to this strip, and when it is tripped the furnace turns on or off. The standard interpretation of this mechanism is that (i) the states of the strip carry information about the temperature of the room, and (ii) this is *why* that component is included in the overall mechanism. Supposing, then, that the thermostat has a range between 50 and 90 degrees, and that the strip is reliable to one degree, there are 40 possible states the strip could be in. If these states are equiprobable, then the amount of information carried on average and by any particular state is approximately $-\log_2(1/40) = 5.322$ bits.

The gun and thermostat are discussed below, but to give an idea of how IT is deployed in cognitive science, consider Linsker's classic (1988) analysis of learning in Hebbian networks. Beginning with a description of how activation propagates through such a network and how Hebbian learning influences connection weights (i.e., a C-story), Linsker constructs an IT-story explaining what it is that the learning algorithm does: It maximizes mutual information between input and output layers, a result known as the 'infomax principle'. So, to take a contrived example, suppose a simple network has two input nodes, and the inputs are [0 1] and [1 0]. Furthermore, these inputs are equiprobable, and the probability of inputs [0 0] and [1 1] are zero. It is certainly possible for the network to map both inputs to the same output, e.g., [1 1]. However, such a mapping looses a bit of information (because the input carries on average one bit of information, but the output carries zero information: $-\log_2(1) = 0$). The infomax principle thus tells us that the network will not converge towards this hypothetical solution, but rather towards outputs that retain the information in the inputs, e.g., [1 1] and [0 0]. Furthermore, notice that as the amount of noise present in the network increases, the amount of mutual information possible between input and output decreases; consequently, the infomax principle predicts that noise-free networks will converge to localist output representations (since they maximize separation between outputs), while those with more noise will exhibit distributed or superimposed outputs. Finally, Linsker also shows that the infomax principle requires that hidden nodes develop causal sensitivities tuned to specific features in the input corpus, i.e., they come to play feature-detection roles. This is clearly an epistemologically rich set of results, and illustrates why IT is so important to cognitive modeling.

These three examples – the gun, thermostat, and the infomax principle – are sufficient to illustrate several 'industry-standard' aspects of information theory, all of which are present above:

(1) IT does not involve any obviously semantic properties. (Floridi 2004, Godfrey-Smith 2007).

(2) C-descriptions and IT-descriptions provide complementary perspectives on the same causal processes.

(3) Information in the IT sense is ubiquitous: It can be applied to any system that exhibits relations between states of components.[5]

To this list we can also add some less-frequently acknowledged features of IT. For instance, a corollary of (2) is:

5 For instance, in addition to the gun, standard philosophical examples of systems that carry information but do not have the function of doing so include footprints, smoke, and tree rings.

(4) IT-descriptions provide insights that C-stories do not.

This observation is clearly illustrated by Linsker's infomax principle, but it is also present in our other examples. Suppose, for instance, we were interested in replacing the bi-metallic strip with some other mechanism, e.g., one constructed from guns. Despite the fact guns are not information-processors, our IT-analyses tell us we need at least six guns (because $2^5 < 40$ and $2^6 > 40$). Alternatively, suppose we have a black box that controls the furnace. We know from observation of its inputs and outputs that it is sensitive down to one-degree increments, and that it only works in the range of 50-90 degrees Fahrenheit. Furthermore, while it is functioning loud bangs are produced, projectiles fly into the room, and there is the smell of burnt gunpowder, i.e., we have reason to believe that the black box uses guns to control the furnace. With our IT-stories in place, we can predict that there will be at least six guns inside the black box. Finally, as a third perspective, suppose that our C-story of the thermostat includes as a continuous variable the angle of the tangent to the curve at the endpoint of the bi-metallic strip. We note that this variable can be in an infinite number of states, so we infer that if we want to replace the strip with some other mechanism we need one that takes on the same number of states. The IT-description saves us from this foolishness by telling us that the causal story does not capture what is relevant about the states of the strip with respect to its capacity to cause the switch to trip in response to the ambient room temperature. In these ways, the perspective provided by IT has epistemological utility distinct from what we get from the causal story alone.

Another observation motivated by our examples is:

(5) The use of IT is guided by functional or teleological attributions.

As noted above, it is commonplace to distinguish between devices or components that have the 'job' of carrying or processing information (e.g., bimetallic strips in mechanical thermostats, hidden units in Hebbian networks) and those that do not (e.g. firing pins in guns, tree rings, footprints, smoke). This distinction is, of course, not unique to IT, but rather is a manifestation of the commonsense principle that mathematical descriptions of a target system are motivated by the goal of describing what a system does and how it does it. For instance, suppose there is an object on the floor keeping a door from swinging shut. We could construct a C-story involving the mass of the object, its geometry, the frictional properties of the surfaces involved, the mass of the door and subsequent force applied to the object, and so forth, the result being an explanation of how the object keeps the door from closing. However, it turns out that holding the door open is an accident; what the object is *actually* doing is disrupting the flow of cool air emitting from a nearby vent (e.g., to facilitate dispersion). As a result, while our C-story is *indirectly* useful insofar as it helps us build doorstops out of similar materials, it is superfluous as a direct explanation since it models the wrong behavior; so we go back to our mathematical toolbox and deploy some other tools in an attempt to understand what the object is actually doing in that context. The point is that in generating C-stories, the tools we use and the behavior we characterize using those tools is guided by beliefs concerning the role of the target system in the context in which it appears. The same holds for IT: Our IT analysis of the gun is indirectly useful (for building a gun-based thermostat), while the analyses of thermostat strips and neural networks presumably describe what those devices have the function of doing.[6]

6 The claim is intended to be neutral with respect to the metaphysics of functions; all that is required for present purposes is that teleological considerations are part of the normative practices governing model construction, and these practices require that acceptable models describe the causal or information-theoretic contributions components make to the relevant behavior of the system (Craver 2007). We see such appeals in the literature; for example, Linsker justifies his

A final observation is implicit in our examples, namely,

(6) IT analyses are often motivated through representation-laden language.

An obvious example of representation-laden language is the standard (and “unfortunate” (Floridi, 2004)) vocabulary of information theory – e.g., ‘surprisal’, ‘expectation’, and ‘uncertainty’ – which (incorrectly) suggest that semantic properties are intrinsic to information theory. Yet we also find other forms of representation-talk in IT contexts. For example, in constructing an IT analysis of RNA encoding, Schneider et. al. (1986) frame the problem as one of quantifying the “information content” of receptor sites. Similarly, Linsker appeals to representation in deriving the infomax principle, e.g., claiming that “input signals are represented in [a] layer [of nodes]” (1988, p. 114). Even if information theory is innocuous regarding semantic properties (observation (1)), it is common to find representation-laden vocabulary used for 'auxiliary' purposes, e.g., motivating or explaining the information-theoretic analysis in non-mathematical terms.

3. … And How to Reject It

With these observations in hand it is possible to appreciate how an attack on indicator representation may or may not impact the use of IT. Recall that indicator representation piggybacks on information theory by 'promoting' states of components that (i) carry information and (ii) have the function of doing so to the status of representations. The basic story is well known (Dretske 1988): If a component has the function of carrying information about some condition, it can *malfunction* in the sense that downstream elements causally respond as if nothing is awry. The component thus *misrepresents* that condition when it malfunctions, and if it misrepresents when it fails, it *represents* when it does not. In this way we get an approximation of truth conditions from standard information-theoretic practice, i.e., IT plus proper function.

Critiques of indicator representation fall into two groups: metaphysical and epistemological. An example of the former is the 'insufficiency argument', which simply points out that it is logically or metaphysically possible for states to satisfy the specified conditions without thereby instantiating semantic features. This strategy leaves IT intact, since indicator representation depends on IT and not vice versa. If the objection has any consequences for how we use IT, they will concern observation (6) – the use of representation-laden auxiliary vocabulary – since, according to the objection, the states so identified have no semantic features. However, since that vocabulary is not being used for its semantic aspects – the analysis being motivated is, after all, non-semantic (observation (1)) – a reasonable response is to simply note through footnote, subscript, or some other method, that the terms are not being used in their semantically-loaded senses.[7] In other words, the insufficiency argument is a paradigm case of a 'merely verbal' disagreement, at least from the perspective of information-theoretic modeling.

A second example of a metaphysical critique of indicator representation is to argue that at least some cognitive processes we assume carry information do not in fact do so; consequently, these

IT analysis by explicitly appealing to the presupposition that the target system has the function of processing information, writing, “Why does a feature-analyzing function emerge from these [Hebbian] development rules? Is it a mere accident or curiosity? Or are the development rules perhaps acting to optimize some quantity that is important to the information processing function of a perceptual system?” (p. 105)

7 For instance, Schneider et. al. (1986) could respond to the insufficiency argument by including a qualification: By 'information content' they mean only 'the probabilistically-characterized causal influence that RNA bindings have on downstream protein expression'. A similar strategy is available to Linsker.

processes cannot use indicator representations. Obviously, this strategy impacts the deployment of IT in cognitive science insofar as it asserts that some IT-stories describe features of the world that simply do not exist, and if we want our theoretical analyses to respect the way the world actually is, we ought not to use IT in such cases. However, given the widespread acceptance that information is ubiquitous (observation (3)), as well as the fact this strategy does not appear in our case studies (section 4), it is bracketed for the remainder of this paper.

A third metaphysical strategy is to allow that our targets carry information, but deny that they have the function of doing so; call this the 'wrong-function' argument. For instance, a detractor might grant that the hidden units of a neural network carry information, but argue that doing so is not their job – instead, hidden units are more like tree rings, firing pins, or footprints in the snow: They participate in causal relations (and hence carry information), but they don't play an information-carrying role, and since they do not have the function of carrying information, they do not represent. This approach to rejecting indicator representation potentially impacts IT by virtue of challenging the intuitions noted in observation (5): If hidden units (for example) do not have the function of carrying information, then our IT-analysis of their behavior (e.g., the infomax principle) does not describe what they are doing – the model is modeling irrelevant behavior. Consequently, the wrong-function argument can have significant implications for standard practice in cognitive science.

In contrast to metaphysical objections, the concern underlying epistemological critiques is that the notion of indicator representation adds nothing to our understanding of the behavior of a target system, even if states of that system *do* instantiate such representations (Chemero, 2009). Unfortunately, the terrain here is a bit complicated. Let's distinguish between two types of epistemic utility: Absolute and relative. Appeals to indicator representation have *absolute utility* if they are epistemically useful parts of the complete explanation of the target behavior, while they have *relative utility* if they are epistemically useful in the construction of a complete explanation, even if they do not appear in the final product.

What impact might denying the absolute utility of indicator representation have on information-theoretic modeling? The answer will depend on the scope of the denial: On the one hand, the putatively superfluous components of indicator representation might be limited to semantic features, leaving the underlying basis for positing those features (i.e., information theory) intact; call this a 'weak' rejection. On the other hand, a 'strong' rejection holds that not only are semantic features superfluous, but so is their basis, including IT. Notice that a weak rejection of absolute utility is arguably a 'merely verbal' attack relative to IT if it can be shown that when they use representation-talk, IT-stories do not presuppose indicator representation even if they presuppose that components can have the function of carrying information; a strong rejection does not leave any such wiggle room.

What about denials of relative utility? As I argue below, representation-talk often plays a heuristic role in the construction of C-stories and IT-stories. In this case, rejecting indicator representation could indirectly impact IT because in doing so we eliminate a tool for generating the IT-story in the first place. Alternatively, it may be that this representation-talk is semantically innocuous – what is actually doing the heuristic work is commonsense teleology and naive information theory, in which case a rejection of representation-talk seems to constitute a 'merely verbal' attack.

4. Yes, No, Maybe So

In the preceding section I considered three examples of information-theoretic analysis, and discussed how different strategies for rejecting indicator representation may (or may not) conflict with the use of information theory. Here I survey three recent critiques of indicator representation with the goal of determining where they fall in this spectrum of possibilities. For each I first indicate why they are properly viewed as 'revisionist' (in the sense of potentially impacting the use of information theory), and then take a critical look at the arguments behind their positions in an attempt to determine whether their revisionism extends to information theory, i.e., whether their revisionism is 'merely verbal' from the IT perspective.

4.1 Ramsey (2007)

Ramsey (2007) is not kind to indicator representation (i.e., the 'receptor notion'). For example, he writes,

> The family of representational notions I'll explore is one I will simply call the “receptor notion.” In the neurosciences, the same sort of state is often referred to as a “detector.” [M]y aim will be to argue that it is not a useful theoretical posit of cognitive science. (p. 118)

And,

> This style of representation [i.e., the 'receptor notion'] often borrows from Shannon and Weaver's theory of information, and rests on the idea that … states represent certain stimuli because of a co-variance or nomic dependency relation with those stimuli. … I argue that the notion is too weak to have any real explanatory value. (p. xvii)

Besides expression a revisionist sentiment, these passages suggest that Ramsey's argument is epistemological; however, a closer inspection reveals that the underlying critique is metaphysical. In broad strokes, his argument has two stages. First, articulate a necessary condition on any state being usefully identified as a representation, and second, argue that indicators do not satisfy this condition. The necessary condition is the “job description challenge”: For an appeal to representation to be epistemologically useful, the putatively representational state must play a “recognizably representational” role (cf., p. 124), where these roles include “informing, denoting, or standing for something else” (p. 25, italics removed). The idea seems straightforward: If a state does not have the function of representing, then calling it a representation is misleading.

The second stage of Ramsey's argument is to claim that the receptor notion fails to rise to the occasion. Bracketing the details[8], it appears that what is being offered is a version of the insufficiency argument. For instance, at one point Ramsey concludes,

> the receptor notion … does not provide us with any sense of how a state or structure actually plays a representational role. The actual stated functional role – reliably responding to some

8 There are many; here is a selective summary. Ramsey identifies three 'recognizably representational' roles: (1) When a state is part of a containing structure that taken as a whole is isomorphic to some domain and downstream components use this isomorphism to guide behavior relative to that domain (e.g., p. 78, p. 201); (2) When a state is a symbolic data structure resulting from the 'encoding' of information (e.g., pp. 68-74); and (3) When a state carries information that is used by a mind to make inferences or form beliefs (e.g., p. 200). Roles (1) and (2) are 'standing in' roles – the isomorphic structure (e.g., a map) stands in for its domain (e.g., a road), and internal, encoded symbolic structures stand in the states they encode (e.g., numerals represent integers). Role (3) is essentially inferential: “The notion of representation is *not* one of standing *in* for some thing or condition, but *implying* or *entailing* some condition. The inference is what converts this mere entailment relation … into a representational relation.” (pp. 200-201; cf. p. 141). Ramsey argues that indicator representations fail to play either of these roles, so to call them 'representations' is a 'conceptual confusion'.

> external conditions – is not, by itself, a role *sufficiently representational* in nature to justify treating some state or structure as a representation (p. 124-5, emphasis added).[9]

The claim, then, is that having the role of carrying information is not sufficient for a state's being representational; therefore we ought to avoid treating it as such. As noted in section 2, the insufficiency argument does not challenge the use of IT; it is 'merely verbal'.

That being said, elsewhere Ramsey shifts from attacking sufficiency to attacking the supposition that states have the function of carrying information. Specifically, he targets the claim, attributed to Dretske, that "if the relevant nomic dependencies are relevant to the proper functioning of the structure, then whatever information is carried by those dependencies must be relevant as well." (pp. 132-3). Ramsey's goal is to show this conditional to be false by providing examples of systems that have relevant nomic relations between states yet their behavior is not illuminated by bringing information theory to bear on those relations. In this way he can show that not only are the conditions for being an indicator insufficient for representation, but also that those conditions are often not satisfied in the first place.

Ramsey's argument proceeds through examples, only some of which can be covered here. The first is a device widely agreed to not have the function of processing information: a gun. Ramsey notes, "The firing pin in a gun functions as a reliable responder to other conditions (i.e., a pulled trigger). Yet, it is clearly not employed, in any serious way, as an information bearer or representational device." (pp. 138-9). As noted in section 2, there is a C-story for the gun. Furthermore, it is certainly the case that the firing pin is designed the way it is (e.g., made of steel rather than lead) because of a need to guarantee those causal relations hold, for otherwise the device would not function properly. We also noted that we can take an information-theoretic perspective on these causal relations, but, since guns are not information processors, the IT-story does not illuminate what the components of the gun do, at least not "in any serious way". In other words, a gun is an example of a device for which components are recruited because of their nomic relations, these nomic relations are the same relations as described by the IT-story, but the IT-story is not relevant to understanding the device, because the components do not have the job of carrying information. Therefore, it is not always the case that the relevance of the nomic dependencies implies the relevance of the information carried by those dependencies, and Dretske's claim is rendered false.

Ramsey takes the point further by suggesting that many devices commonly viewed as having the function of processing information do not in fact have that function; rather, they are more like guns. For instance, our beloved thermostat gets the following treatment:

> If we look closely at the functional architecture of the thermostat … it is far from clear why we should say that information carrying is the functional role of the strip, as opposed to reliably responding to certain conditions. (p. 136).

As in the case of the gun, the goal is to show that components can be recruited into a system for their nomic dependencies (a safe assumption), those relations can carry information (trivially true on account of the ubiquity of information), but this does not imply that the function of those components is to carry information. In this case, the claim is that the bimetallic strip of the thermostat is just such a

9 In this passage Ramsey is criticizing a non-teleological account of indicator representation. However, the same claim is arguably made regarding full-blooded indicator representation. For instance, he writes, "The problem with receptor representations is not that the systems that allegedly use them *can* be given a purely causal … non-representational characterization. Rather, the problem is that the theoretical frameworks that invoke them … actually assign to them a non-representational role." (p. 142)

component. Consequently,

> the problem [with Dretske's account] is that a structure can be employed in such a way that the causal and nomic relations that enable it to carry information *are* explanatorily relevant while the information resulting from such relations is not. (pp. 137-8).

If the function of the bimetallic strip is *not* to carry information about the temperature so that the switch can 'use' that information to 'decide' whether to trip the furnace, then the IT-description is misapplied, and at best is indirectly useful (like our IT description of the gun).[10]

Finally, Ramsey argues that the same argument applies to artificial neural networks. For instance, he considers whether "clusters in vector space reveal that the receptor-like hidden units are serving to encode distinct chunks of information about the world" (p. 145), and rejects the proposal:

> To draw the conclusion that a vector analysis reveals a network's representation scheme, you have to first assume the hidden unit vectors are, in fact, serving as representations. [But] The mere clustering in the response profile doesn't show this. There are, after all, plenty of non-representational systems that admit of the same sort of vector clustering.

The issue for network states is the same as that for guns, thermostats, and his other examples: Hidden units causally respond in to certain stimuli, but this does not mean that they have the *function* of carrying information, as illustrated by the fact that other sorts systems causally respond in similar ways and clearly do not have such a function. Consequently, while we could construct an IT description of a network, it wouldn't tell us anything about what the network is actually doing; to understand what the network is actually doing, all we need is the C-explanation. Information-theoretic descriptions – such as the infomax principle – are perhaps indirectly interesting, but they do not add anything to our understanding of the behavior of the network itself.[11]

To sum, Ramsey's rejection of indicator semantics appears to involve both the insufficiency argument and the wrong-function argument. While the former leaves the motivation for applying information theory in place at the cost of qualifying terminology, the latter removes the motivation for and (direct) epistemological payoff of applying information theory by insisting that it is being misapplied to systems that do not have the job of processing information. Consequently, Ramsey's revisionist argument is far from 'merely verbal' – it is genuinely radical.

4.2 Hutto & Myin (2013)

In what amounts to a precis of Hutto & Myin (2013), Hutto (2013) writes:

> Proponents of [radical enactive cognition] argue that the standard characterizations of cognition

10 This strategy is also used as rejoinder to those who respond to Van Gelder's (1995) infamous 'Watt governor argument' by asserting that the governor's arms carry information about the speed of the engine (i.e., Bechtel 1998). According to Ramsey, the arms of the governor are no different than firing pins or bimetallic strips: "There really is *no* interesting sense in which the weighted arms play the functional role of carrying information about something else [e.g., the speed of the engine]" (p. 213). Consequently, C-stories are useful for understanding how a centrifugal governor works, but IT-stories are not.

11 Apropos the present point, Ramsey concludes that "the receptor notion can undermine scientific progress by promoting a misguided analysis of the functional role of internal cognitive states" (p. 148). Again, the claim is not that having the function of carrying information is insufficient for representation, but rather that cognitive states do not have that functional role.

> and cognitive processes, in terms of representation and computation, are misguided. Representation and computation do not lie at the heart of all cognition and do not form its basis. (Hutto, 2013, p. 142).

A parallel claim from Hutto & Myin (2013; H&M) is:

> Cognitive systems don't "pick up" or "take in" any informational contents; there are no such things as informational contents to take in. Such imagined contents are not "objective commodities," and cognitive systems do not "traffic" in them... (p. xvi).

These assertions raise the issue dealt with in this paper: If the notion that cognitive systems process information or represent is rejected, then what is the status of information theory?

Taking their cue from Dretske (1988, p. vii, quoted in H&M, pp. 65-66), they conceive of informational content as a "commodity" that is 'carried' or 'transmitted' by causal processes. As a form of content, it has semantic features, such as truth conditions or reference (p. 67). H&M's argument then takes the form of a dilemma: Either this peculiar commodity is already out there, floating around in the world, free to be used in our theories of representation, or there must be a convincing story about how it comes into existence from 'merely causal' processes. The first option is not open to the naturalist, and furthermore, an 'explanation' of representational content in terms of preexisting representational content isn't much of an explanation. The second option – "the hard problem of content" (p. 69) – is also untenable because, H&M claim, indicator semantics fails to bridge the gap between mere causal processes and genuine semantic properties.

Unfortunately, H&M are not clear about *why* indicator semantics fails to bridge that gap. Regardless, their position appears to involve a version of the insufficiency argument. For instance, the name they choose for their challenge invokes Chalmers' (1995) infamous 'hard problem of consciousness', and the core claim of that argument is that physical descriptions alone are not sufficient to imply consciousness, suggesting that H&M are making a parallel insufficiency claim about the proposed conditions on being an indicator representation. Similarly, H&M cite Ramsey (2007) with approval, summarizing the relevant point as being that states that have only the properties of indicator representations "do not automatically qualify as truly contentful, thus representational" (p. 62). Again, this is an instance of the insufficiency argument. Consequently, their critique appears to be 'merely verbal' from the perspective of IT.

More importantly, H&M acknowledge that information theory is not impacted by their position.[12] For instance, concerning their dilemma they write,

> One might opt to be impaled on the first horn … and retain the scientifically respectable notion of information-as-covariance. This would allow one to retain full naturalistic credentials while relinquishing the idea there is such a thing as informational content – the path we recommend (p. 68)

By 'information-as-covariance' H&M mean information-theoretic information. So, for instance, they distinguish between that notion and 'information-as-content', where the latter but not the former has semantic properties (p. 67).[13] In short, in contrast to Ramsey, H&M are relatively clear regarding the

12 Indeed, it sometimes appears that their version of the insufficiency argument consists in nothing more than the standard observation that information theory is 'syntactic' and not semantic. See, for instance, the multiple attempts at stating the argument on pp. 67-68.

13 H&M also reiterate their claim is insufficiency, writing, "covariation in and of itself neither suffices for nor otherwise

intended impact of their representational revisionist critique on current practice in information-theoretic modeling: Stay calm and carry on. As I suggested in section 2, this may involve explicitly or implicitly qualifying the use of representation-talk in that type of research, but from that perspective the issue is purely verbal. So, in response to the main question posed by this paper, 'no', H&M's revisionist position does not impact the use of information theory in the cognitive sciences.

4.3 Chemero (2009)

Chemero (2009) explicitly takes issue with the epistemology of indicator representation rather than its metaphysics. Specifically, he allows that components may have the function of carrying information, and that doing so is sufficient to instantiate indicator representations, but denies that explanations appealing to these representations add anything to our understanding of the systems they describe, at least if we already have a dynamical systems theory (DST) explanation in place. For instance, he writes,

> For these [DST] models … the teleological, representational story doesn't seem any more informative than saying that the robots [modeled by the DST models] evolved or were designed for their tasks. (p. 71)

And,

> the representational story could be told but … that's not particularly relevant, because the dynamical systems theory explanation tells us everything important about the system. (p. 72)

The issue, then, is whether Chemero's charge of epistemological futility applies to IT-stories as well as the representational accounts of which they are a part.

In what follows, I consider two of Chemero's epistemological arguments – the *dependency argument* and the *predictive argument* – and show that the former amounts to a strong rejection of absolute utility, while the latter can be treated as either strong or weak as regards absolute utility. Furthermore, since both of these arguments admit of fairly straightforward objections, I consider how their implications for information theory change when adjusted to accommodate these objections. Specifically, I map the complex terrain that must be navigated if we are to fully assess how epistemological critiques relate to IT, leaving the task of finding a path through this terrain to those who advocate epistemological revisionism.

The dependency and predictive arguments appear in Chemero's discussion of Husbands, Harvey & Cliff's (1995) neural-network controlled mobile robot, for which there is both a DST explanation (provided by the authors) and a representational account (proposed by Chemero for purposes of illustration):

> There are reasons to prefer the dynamical account to the representational one … First, the representational story depends on the dynamical story about the control system [of a robot], but not vice versa. … [i.e.,] the representational description is dependent upon the dynamical one. [Second,] the representational description of the system does not add much to our understanding of the system. Once we have the full dynamical story, we can predict the behavior of the robot in its environment completely, and we can do so without making reference to the representational content of any states of its control system. … [i.e.,] the representational gloss

constitutes or confers content" (p. 67)

> does not predict anything about the system's behavior that could not be predicted by the dynamical system alone. (p. 77)

As suggested by this passage, the *dependency argument* goes something like this: If a representational explanation depends on a prior DST explanation, then the former is superfluous. Representational explanations depend on a prior DST explanation, so representational explanations are superfluous.

The first premise of the dependency argument appears to be an instance of a more general principle, namely that any putative explanation that is constructed on the basis of a prior explanation is superfluous. If this is correct, then the dependency argument is a strong critique, since IT-descriptions are constructed on the basis of prior causal models. For instance, Linsker's infomax principle is only arrived at after the causal behavior of Hebbian networks is fully described; similarly, Williams & Beer's (2010) IT-description of the control network for a simulated robot is only developed after their DST analysis is in place. Finally, we should note that even the IT-story for the gun (section 2) was grounded in a prior conception of how the components causally interact. Therefore, if the dependency argument is sound, IT-stories are epistemically superfluous.

Also on display in the passage quoted above, the *predictive argument* holds that since DST explanations fully predict the behavior of a target system, representational explanations are superfluous.[14] This is not to say that appeals to representation are never predictive – after all, if we have some idea of the representational role a component plays, we can make inferences about the effects that component has on other parts of the containing mechanism. (For example, cells in visual cortex may have the function of detecting edges of a certain orientation in the visual field, so we can predict, among other things, that downstream components tasked with object-recognition will be causally sensitive to ensembles of such feature detectors.) So a suitable interpretation of the argument is that once we have a fully-developed mathematical account of the causal behavior of a system, any prior representational description can be discarded, because the C-story takes care of all our prediction needs.

The predictive argument admits of both weak and strong interpretations, depending on what we are throwing away when we discard a representational description. On the one hand, the claim may be that it is the semantic features of indicator representations alone that ultimately fail to provide predictive leverage when a DST description is also available, i.e., knowing the information processing function of a component does help predict behavior, but contents, truth values, and the like do not. In this case, the predictive argument is weak with respect to absolute relevance. On the other hand, the claim may be inclusive of information theory (or teleology in general), in which case the predictive argument is strong.

What we find, then, is that when taken at face value, one of Chemero's epistemological arguments (the dependency argument) directly challenges IT, while the other (the predictive argument) may or may not directly challenge IT. However, there are some straightforward objections to Chemero's arguments, and accommodating these may change our assessment. So, in the remainder of this section I will critically adjust those arguments.

Regarding the dependency argument, Chemero writes, "one *must* have the dynamical story first, before one can concoct a representational story" (p. 73, emphasis added). However, it is possible to challenge this claim, because indicator representations appear to play a role in the construction of causal (or DST) explanations. Consider, for instance, the work Chemero appeals to in motivating his argument: Harvey, Husbands, and Cliff's (1995; henceforth HHC) DST analysis of an evolved neural-

14 The predictive argument comes up several times in addition to the passage quoted above: "Why should we bother with representational explanations when we have precise, perfectly general, counterfactual-supporting mathematical ones?" (pp. 72-73). And, "If one has the complete dynamical story, what is left to be explained?" (p. 73).

network controller. Chemero's claim is that HHC's DST is not the result of appeals to indicator representation, but a closer look suggests otherwise.

HHC's (simulated) holonomic robot is tasked with navigating to the center of a circular arena. The robot has tactile and visual sensors that detect collisions and ambient light. The controller is an artificial neural network with a fixed number of sensor and actuator nodes, but the number of hidden nodes and their connections can vary. The controller is constructed indirectly using a genetic algorithm, and, having evolved a controller that successfully solves the task, HHC proceed to analyze the controller using the tools of dynamical systems theory. However, contra Chemero's claim, to facilitate the DST analysis HHC simplify the network by drawing on the notion of indicator representation. Specifically, they consider the correlation between type of information being presented – visual or tactile – and the responses of hidden nodes, and remove those nodes "that play no part in visually guided behaviors" (1995, pp. 95-6). In this way, HHC use patterns of causal co-variation to focus their analysis only on the sub-network that has the function of processing visual information, i.e., on the sub-network that represents visual stimuli.

Appealing to indicator content to aid in arriving at a DST-story analysis also appears in other cases. For instance, Williams & Beer (2010; see also Beer 2003) give a dynamical analysis of an evolved neural network controller for a simulated robot restricted to move horizontally at the bottom of a two-dimensional, rectangular arena (i.e., the bottom of the screen). The robot is presented with a circle, which it can sense using 'rays' projected from its sensors. A second circle then descends from above, and the robot is tasked with catching (i.e., intersecting the path of) this object if it is smaller than the first and avoiding it if it is not. This task requires that the robot compare the second circle with the first, so it stands to reason that some components of the network have the job of storing the size of the first circle so that it can be compared with that of the second; a dynamical analysis of the network will articulate how this is achieved. Consequently, to construct a DST explanation we need to isolate the components that perform this function so that their causal contributions can be described:

> To determine this, we examples the state variables of the agent – its neural outputs and body position – at the time when the first object is removed, and found that a single neuron … correlates with the size of the object at this time. (pp. 39-40)

Like HHC, Williams & Beer appeal to the fact that a sub-network of the controller has the function of carrying a certain type of information (in this case the size of the first circle) so that they can develop a DST explanation capturing the relevant causal processes, i.e., those that explain how the robot does what it does. Far from being dependent upon a prior DST explanation, appeals to indicator representation are used to guide the construction of dynamical analyses.

Examples such as these show that it is hard to treat one type of description (causal, DST) as prior to another (representational), as required by the dependency argument. Instead, representation-talk plays a heuristic role, guiding the construction of causal or DST explanations, and causal or DST explanations help refine representational ascriptions (Neander 2006); the result is a sort of 'epistemological feedback loop'. Consequently, the dependency argument is based on a mistaken understanding of the epistemological role of indicator representation.

For the sake of argument, let's suppose that despite appearances, these appeals do not actually contribute to the construction of the final causal, DST, or IT stories. They are slips of the tongue, shorthand for some nonrepresentational approach to model construction, or unfortunate methodological lapses. In this case we once again have an issue concerning scope of denial. First, the revisionist might simply be urging us to stop using words such as 'representation' or 'content' in these heuristic contexts, because these uses don't actually appeal to contents or other semantic features; instead, they rely only on (presumably unproblematic) teleological, causal, and information-carrying assumptions. On this interpretation, the critique is weak with respect to relative utility (because it preserves the basis of indicator representation), and hence appears to be 'merely verbal'.[15] Second, the revisionist might be claiming that these are the wrong heuristics entirely, e.g., that in constructing a causal, DST, or IT story we either *ought* not to be appealing to components having the function of carrying information (and hence not to representation), or we *are* not making such appeals, even though it seems like we are. This claim is much harder to defend, for the revisionist must either provide an equally effective alternative toolset that will aid in the construction of models, or has to explain away the appearance of appealing to the function of carrying information in constructing DST models. Either way, the scope of denial is more expansive, and as a result it is much harder to see how an IT story (or a DST story, for that matter) could be be constructed; consequently, the critique is strong with respect to its impact on information-theoretic modeling.

To sum, the dependency argument asserts that indicator representation does not enter the scene until after a causal or DST story is constructed, but there is plenty of evidence that this claim is false: Appeals to indicator representation play a heuristic role by directing the construction of causal, DST and IT stories. The question then becomes what impact the denial of this role might have on the use of IT, and, once again, the answer to this question depends on the scope of that denial: If IT and naïve teleology is retained, the critique is weak and concerns merely the improper use of representational vocabulary, but if IT and naïve teleology is rejected along with representation-talk, it may be difficult to construct IT stories at all. This shows that epistemological critiques of indicator representation *might* be more than 'merely verbal' - they could be recommending fundamental changes to scientific methodology.

What about the predictive argument? First, the preceding discussion shows that the predictive capacities of representation-talk may not always be superfluous, since representation-based inferences arguably play a role in DST-story construction by guiding our attention to certain causal trajectories to the exclusion of others. Consequently, one branch of investigation leads back to same terrain covered by our earlier discussion of the dependency argument. Second, Chemero seems to assume that prediction is the only type of epistemic utility we care about. However, teleological explanations (of which indicator representation is a species) are traditionally understood as answering a different type of 'why' question than causal stories (Canfield, 1964): Rather than tell us the sequence of causal events that lead to a present or future event, representational explanations give reasons for why a particular component is present in a containing system. For example, in his influential discussion of indicator representation, Dretske (1988, p. 84) notes that the representational function is "an explanatorily relevant fact about [the component] itself", not about what it causes. More generally, indicator representation may have absolute utility even if that utility is not related to the behavioral predictions it affords; as far as I can tell, Chemero does not discuss the literature on teleological explanation. Consequently, a comprehensive investigation into the consequences of epistemological revisionism for information theory will require substantive elucidation in this regard. For instance, it may be that the revisionist claims that all teleological explanation is superfluous, because these alternative 'why' questions don't really add to our understanding, or can be translated into causal 'why' questions, or for some other reason. Alternatively, perhaps it is merely representational explanations that fail to involve teleology (and hence do not qualify as teleological explanations, contra Dretske's assumption). In that case we would still be allowed to talk about components having the function of carrying information. In short, how the epistemological critique affects information-theoretic modeling may depend on how that critique treats teleological explanation in general.

15 Furthermore, this interpretation preserves the original dependency argument insofar as it is possible to assert that the 'real' representational stories arise only after a C-story (i.e., a DST-story) has been fully constructed.

To sum, the predictive argument seems to overlook the heuristic role that representation-talk can play as well as philosophical traditions that grant teleological explanations an epistemic status distinct from that of causal stories. As with the dependency argument, the relevance of epistemological critiques of indicator representation for information-theoretic modeling turns on how the resulting 'conceptual tangles' are dealt with by the proponent of such critiques. Consequently, for present purposes, our answer to the question driving this paper is 'maybe' - it is unclear whether epistemological critiques of indicator representation will have consequences for the use of information theory.

4. Conclusion

While information theory is generally viewed as philosophically unproblematic, theories of mental representation built upon it are not: Representational revisionist philosophers increasingly call on cognitive science to divest itself of such 'indicator' or 'receptor' notions of representation. Given the importance of IT to cognitive science it is reasonable to inquire into the status of information theory under these negative critiques. I've argued that current revisionist arguments appear to have the bases covered: (1) Ramsey's (2007) argument implies that the use of information theory should be curtailed on the grounds that at least some devices presumed to process information in fact do not have that function; (2) Hutto & Myin's (2013) approach makes room for information theory sans content; and (3) the fate of IT under Chemero's (2009) approach depends on how the heuristic and teleological epistemological roles of indicator representation are accommodated. These results also show that revisionist arguments, even when limited to the context of information theory, need not be 'merely verbal'. Instead, some revisionist arguments have radical consequences for how we understand cognitive processes, i.e., not only sans representation, but also *sans information*.

One might object to my focus on information-theoretic modeling; surely, even if revisionist arguments do not bear on the use of information theory, they may have consequences for other aspects of cognitive science. This is correct, but misses the more general point of the paper. The reason I chose to focus on information theory was that it is part of the basis for indicator representation; but we could have easily discussed the relation between revisionist arguments and other types of epistemic endeavors, e.g., computational modeling, mathematical modeling, or experimental design. For example, what guidance does rejecting indicator representation provide for choosing the algorithms and data structures used in a computational model, or for choosing between competing sets of equations?

This is the fundamental challenge facing representational revisionism: It seems to offer no guidance for these types of endeavors, and hence is merely verbal. On the one hand, if the response is that it *does* offer guidance, then proponents of revisionism should articulate not only what that guidance is, but also how the results are different from – and more useful than – the approaches currently available. On the other hand, if the response is that it doesn't offer guidance on those particular endeavors, then we would like proponents to articulate precisely where it *does* offer guidance, what it is, and how it differs (usefully) from current approaches. Either way, we get a better picture of what the contributions of representational revisionism to cognitive science are supposed to be. The present paper is an attempt to fill this lacunae, even if it is limited to the context of information theory.[16]

16 Of course, what we would really like out of a revisionist account is some new algorithm, data structure, set of mathematical equations, etc., of sufficient specificity to allow comparison with currently available tools (e.g., the specification of a new algorithm should allow for complexity analysis). My argument has been limited to simply arguing for consistency or inconsistency with IT, but it goes without saying that nobody is going to stop using information theory just because a philosopher has a bone to pick with representation; instead, if a critique of

representation implies that information theory should be jettisoned, then the critique must be wrong. In other words, the real challenge to representational revisionists is to navigate between irrelevance and reduction to absurdity by showing how the revisionist position leads to novel mathematical, computational, or experimental techniques.

Acknowledgments: Thanks to Tad Zawidzki for comments and discussion on an earlier draft of this paper. Matthew Fricke and John Poland provided valuable feedback during the genesis of the initial draft, as did discussants at the 2011 annual meeting of the Southern Society for Philosophy and Psychology, where a portion of section 3.3 was previously presented.

References

Bechtel, W. (1998). "Representations and cognitive explanations: assessing the dynamicist's challenge in cognitive science." Cognitive Science 22(3): 295-318

Canfield, J. (1964). "Teleological explanation in biology." The British Journal for the Philosophy of Science 14(56): 285-295.

Chalmers, D. J. (1995). "Facing up to the problem of consciousness." Journal of Consciousness Studies 2: 200-219.

Chemero, A. (2009). Radical Embodied Cognitive Science. Cambridge, MA, MIT Press.

Clark, A. (2013). “Whatever next? Predictive brains, situated agents, and the future of cognitive science.” *Behavioral and Brain Sciences* 36: 181-253.

Dimitrov, A. G., Lazar, A. A., Victor, J.D. (2011). “Information theory in neuroscience.” *Journal of Computational Neuroscience* 30(1): 1-5.

Dretske, F. (1981). Knowledge and the Flow of Information. Cambridge, MA, MIT Press.

Dretske, F. (1988). Explaining behavior. Cambridge, MA, MIT Press.

Dreyfus, H. L. (1979). What Computers Can't Do: The Limits of Artificial Intelligence. New York, Harper and Row.

Floridi, L. (2004). Information. The Blackwell Guide to the Philosophy of Computing and Information. L. Floridi. Oxford, Blackwell Publishing: 38-61.

Fodor, J. A. (1990). Psychosemantics or: Where do truth conditions come from? Mind and Cognition: A Reader. W. Lycan. Oxford, Blackwell Publishers: 312-337.

Freeman, W. J. and C. A. Skarda (1990). Representations: Who Needs Them? Brain Organization and Memory: Cells, Systems, and Circuits. J. L. McGaugh, N. Weinberger and G. Lynch: 375-380.

Godfrey-Smith, P. (2007) 'Information in biology', in *The Cambridge Guide to the Philosophy of Biology*, David L. Hull & Michael Ruse (eds.), Cambridge: Cambridge University Press, pp. 103-119.

Husbands, P., I. Harvey, & Cliff, D. (1995). "Circle in the round: state space attractors for evolved sighted robots." Robotics and Autonomous Systems 15: 83-106.

Hutto, D. D. (2013). "Exorcising action oriented representations: ridding cognitive science of its Nazgul." Adaptive Behavior 21(3): 142-150.

Hutto, D. D. and E. Myin (2013). Radicalizing Enactivism: Basic Minds without Content. Cambridge, MA, MIT Press.

Linsker, R. (1988). "Self-organization in a perceptual network." IEEE Computer 21(3): 105-117.

Milenkovic O, Alterovitz G, Battail G, Coleman TP, Hagenauer J, Meyn SP, et al. Introduction to the special issue on information theory in molecular biology and neuroscience. IEEE Transactions on Information Theory. 2010; 56(2):649–652.

Millikan, R. G. (1989). "Biosemantics." Journal of Philosophy 86: 281-297.

Neander, K. (2006). "Content for cognitive science." *Teleosemantics*. G. Macdonald and D. Papineau. Oxford, UK, Oxford University Press: 167-194.

Ramsey, W. (2007). Representation Reconsidered. Cambridge, Cambridge University Press.

Schneider, D., Stromo, G., Gold, L. & Ehrenfeucht, A. (1986) 'Information content of binding sites on nucleotide sequences', Journal of Molecular Biology, 188, 415-431.

Thornton, C. (2013). "A new way of linking information theory with cognitive science." *Cooperative Minds: Social Interaction and Group Dynamics. Proceedings of the 35th Annual Meeting of the Cognitive Science Society*, M. Knauff, M. Pauen, N. Sebanz, I. Wachsmuth (eds.), Austin, TX: 3545-3550.

van Gelder, T. (1995). "What might cognition be, if not computation?" *Journal of Philosophy* 91(7): 345-381.

Williams, P. and R. Beer (2010). Information Dynamics of Evolved Agents. From Animals to Animats 11. S. Doncieux, B. Girard, A. Guillot et al, Springer Berlin Heidelberg. 6226: 38-49.

Williams, P. L. and R. D. Beer (2011). "Generalized measures of information transfer." arXiv:1102.1507v1 [physics.data-an], retrieved 2013-08-27