

INFORMATION AND SUSTAINING MECHANISMS IN FODOR'S THEORY OF CONTENT

Información y mecanismos de sustentación en la teoría fodoriana del contenido

Bernardo Aguilera*

ABSTRACT

According to Fodor's informational approach, mental symbols have content by virtue of standing in certain nomic relations with their referents. These relations are sustained by computational mechanisms which enable the causal route linking mental symbols with the world. Fodor claims, however, that specifying the structure of those sustaining mechanisms is irrelevant for a theory of content. This paper argues that, on the contrary, without an account of the computational constraints under which those mechanisms operate, Fodor's theory is at best incomplete, and incapable of explaining what makes us the only known computing machine capable of bearing mental symbols so far.

Keywords: Informational approach to content, computational theory of mind, mental symbols, reference, asymmetric dependence theory.

* Centro de Bioética, Facultad de Medicina, Clínica Alemana, Universidad del Desarrollo, Departamento de Bioética y Humanidades Médicas, Facultad de Medicina, Universidad de Chile. Santiago, Chile. Correo electrónico: baguilera@udd.cl

Acknowledgments: The author would like to thank Stephen Laurence, Dominic Gregory, Kathy Puddyfoot and anonymous referees for helpful comments on earlier drafts of this paper.

Artículo recibido el 14 de diciembre de 2015. Aceptado el 20 de enero de 2016.

RESUMEN

Según la propuesta informacional del contenido de Fodor, los símbolos mentales poseen contenido en virtud de situarse en cierta relación nómica con sus referentes. Estas relaciones son posibles gracias a la posesión de “mecanismos de sustentación”, que en términos generales consisten en sistemas computacionales que habilitan la ruta causal que conecta los símbolos mentales con el mundo. Fodor afirma, sin embargo, que especificar la estructura de estos mecanismos de sustentación es irrelevante para una teoría del contenido. En este artículo argumento que, por el contrario, la propuesta de Fodor está al menos incompleta si no explica cómo son estos mecanismos computacionales y los límites dentro de los cuales estos operan, y asimismo resulta incapaz de explicar qué nos hace capaces de poseer símbolos mentales.

Palabras clave: Propuesta informacional del contenido, teoría computacional de la mente, símbolos mentales, referencia, teoría de la dependencia asimétrica.

1. INTRODUCTION

According to the dominant paradigm in philosophy of mind, the nature of thought can be explained by appeal to the mechanical manipulation of mental symbols. But while computation theory has provided solid grounds for explaining how symbolic processes could be mechanised, the issue of explaining how symbols could really symbolise something –viz. how they could have content or semantic properties– is still matter of intense debate. One of the most promising accounts of this matter has come from informational approaches to content. They can be regarded as stemming from the empiricist tradition in philosophy, which starts from the rather intuitive premise that we obtain knowledge from the world by getting information about external objects through our senses. Information theory has attempted to refine the notion of information in terms of an objective commodity that can be generated and transmitted. It has developed along with computation theory offering ways of understanding how computing systems such as the mind can pick up and process information about the environment (Dretske 1981; Stalnaker 1984; see Adams 2003 for a review).

Taken at face value, informational approaches claim that mental symbols express the properties of the environment they carry information about. And the way symbols obtain information is by virtue of holding a nomic (i.e. lawful) relation with their referents, in the sense that changes in the nature of symbols can be reliably correlated with changes in their referents. But informational approaches need many refinements to become a convincing view about content. First, they face the risk of falling into a *reductio ad absurdum*, given that information-bearing

states are rife in the natural and artificial world. This is indeed true in the case of simple artefacts. For instance, the mercury column of a thermometer bears information about the local temperature, and the end of a compass needle carries information about where is the north magnetic pole. But surely a theory that implies that thermometers and compasses possess mental symbols is misleading.

Secondly, when considering the case of more sophisticated artefacts such as robots capable of engaging in complex informational relations with the environment, we might wonder why they lack mental symbols¹. Even if informational approaches could avoid the implausible ascription of mental symbols to simple artefacts, they still need to explain how to draw the line between us and other more complex information-processing computing machines –and indeed between us and other biological computing machines that lack mentality (see Aguilera 2015 for discussion). In sum, if the symbols involved in mental processes are supposed to have content because they stand in nomic relations with their referents, it has to be explained what is special about those symbols, such that sets them apart from the information-bearing states of non-minded computing agents capable of processing information from the world.

One of the most influential developments of informational approaches to content has been carried out by Jerry Fodor (1987; 1990; 2008). He is aware of the aforementioned problems of informational accounts, and has centred his efforts on specifying what is special about the informational link with the environment that characterises genuine mental symbols. For present purposes I will focus on his initial discussion in *Psychosemantics*², the main tenets of which have remained almost untouched in his subsequent work. In that book, Fodor puts forward two ways in which symbol-world relationships have to be constrained in order to establish a genuine semantic relation. Those constraints are presented as a means to solve two *prima facie* objections that any informational approach has to deal with. I present them below.

The first objection has to do with the nomic nature of symbol-world relationships. If there is a nomic connection between A and B, it is nomologically necessary that if A is the case, then B. Consider the mental symbol ‘cow’³. Given that it is nomic related to cows, then every time a cow is instantiated in the world a corresponding tokening of ‘cow’ has to be instantiated. But this is not true of ordinary mental symbols insofar as not every instantiation of a cow actually causes a ‘cow’ token. For instance, just a minimal fraction of the cows that exist

¹ Examples of sophisticated artefacts of this sort can be the two robotic rovers that are currently exploring the surface of Mars, called “Opportunity” and “Curiosity”.

² Unless otherwise indicated, all references are to this book.

³ To distinguish a mental symbol from its referent, I shall use quoted formulas to express the former.

in the world happen to cause tokenings of ‘cow’ in my mind, and moreover, cows that exist in isolated places might never be the cause of ‘cow’ tokenings at all. So, what informational approaches have to explain is how it is that just some cows cause ‘cow’, and note that to avoid being question-begging these constraints have to be specified without appeal to other mental symbols. Let us call this worry the ‘not-all problem’.

The second worry has been labelled by Fodor as the ‘disjunction problem’. If the mental symbol ‘A’ refers to A by virtue of being nomically connected with As, then it cannot also be nomologically connected with non-As, such as Bs, since in that case ‘A’ would be referring to the disjunct (A or B). For example, it is plausible to conceive that some ‘cow’ tokens can sometimes be caused by horses by virtue of some nomic relation holding between properties of horses and ‘cow’ tokens, but if this is the case then ‘cow’ would not refer just to cows but to (cows or horses). This is called the disjunction problem since mental symbols are normally supposed to bear reference relations to some particular property and not to a disjunct (and less to an open disjunct as it turns out) of properties as an informational approach appears to imply.

Fodor deals with these two worries by appealing to sustaining mechanisms and his famous asymmetric dependence theory, both of which I will explain in the following section. While the asymmetric dependence theory has received more attention in the literature, in this paper I will challenge Fodor’s proposal concerning sustaining mechanisms. In section 3 I contend that his claim that sustaining mechanisms do not determine the content of mental symbols, ends up leaving unexplained part of what his theory of content is supposed to explain, viz. what is special about the symbol-world relationships possessed by mental agents, such that set them apart from other computing agents that process information from the environment.

2. FODOR’S PROPOSAL

2.1. Sorting the not-all problem: Sustaining mechanisms

Fodor’s way of dealing with the not-all problem is to specify certain sufficient conditions for the instantiation of ‘cow’ tokens such as that when those conditions are met, cows cause ‘cow’. Since those conditions are supposed to be stated in non-symbolic terms, they would allow an informational approach to explain in a non-question-begging way why not all cows actually cause ‘cow’ tokens. Fodor develops a twofold process to account for those conditions. The first corresponds to the encoding of information by transducers, which he describes as

purely psychophysical. It involves the conversion of energy of the environment onto electrical signals that encode information in the brain. For example, there are certain conditions under which red objects cause the tokening of an inner state carrying information about the redness of the object. In Fodor's words:

Psychophysics purports to specify what one might call an 'optimal' point of view with respect to red things; viz., a viewpoint with the peculiar property that any intact observer who occupies it must –nomologically must; must in point of psychophysical law– have 'red there' occur to him (p. 115).

But of course, mental symbols are not just the output of transducers, and so Fodor's proposal needs to accommodate the case of symbols that are generated by inferential processes⁴. Mental symbols of distal environmental objects such as horses or trees cannot be generated just by means of transducers and so their representation cannot be explained by mere appeal to psychophysical law. Psychophysical circumstances can tell you when someone will "see" a horse, but not when she will "see as" a horse. As Fodor claims, "there are no psychophysically specifiable circumstances in which it is nomologically necessary that one sees horses as such" (p. 117).

Here Fodor adds a second step to the process, which corresponds to the mediation of inferences. Mental symbols such as 'horse' or 'tree' are mediated by inferential processes drawn from the perceiver's "background cognitive commitments" (p. 117). The idea is that after transducers encode environmental information, perceptual processes pick up that information and through computation and integration with stored information generate a mental symbol. These mechanisms that support the symbol-world relation correspond to what I am calling sustaining mechanisms⁵. But as Fodor recognizes, to appeal to inferences and background commitments can be question-begging, since they presumably involve previous representations and theories from which we draw the inferences that allow us to token 'horse' or 'tree'. This is clearer when it comes to symbols of more abstract kinds such as protons, which require certain scientific knowledge to be tokened. Fodor avoids the criticism that his account is question-begging by saying that the symbolic capacities involved in sustaining mechanisms are not

⁴ See Fodor (1983) for his account of perceptual systems, where he distinguishes between transducers and input modules. On his view, only the latter are capable to carry out inferences and deliver mental symbols.

⁵ The term "sustaining mechanisms" is not Fodor's, who uses similar expressions to present the same idea, e.g. "mediating mechanisms" or "mechanisms that sustain" (Fodor, 1990: 100). However, since Margolis (1998) it has become a standard way of naming what Fodor intended to characterise as the causal mechanisms that lock a symbol-world relation.

determinants of the content of a mental symbol. As he says, “for the purposes of semantic naturalisation, it’s the existence of a reliable mind/world correlation that counts, not the mechanisms by which that correlation is effected” (p. 122). And since the mechanisms required to sustain the fixation of content are computational, they can, according to Fodor, be specified in causal-syntactic terms, without appeal to symbolic notions. I will return to this issue in section 3 when putting forward a critique to Fodor’s proposal.

2. 2. Sorting the disjunction problem: Asymmetric dependence theory

Concerning the disjunction-problem, Fodor attempts to distinguish between the nature of the nomic relations a mental symbol holds with its referents, and the nomic relations it might establish with anything else distinct from them. So for example, in order to avoid ‘cow’ tokens being about a disjunct such as (cows or horses), there must be some way to tell apart the nomic relation ‘cow’-cows from ‘cow’-horses. Fodor’s suggestion is to state that the causal route between the mental symbol and its referent is special in the sense that it does not depend on any other relation to exist. Hence ‘horse’ tokens are supposed to be caused by cows only because ‘cow’ tokens are, and not vice versa. Put in Fodor’s terminology, the point is that: the causal connection between cows and ‘horse’ tokenings is, as I shall say, asymmetrically dependent upon the causal connection between horses and ‘horse’ tokenings (p. 108).

The same idea can be framed in terms of counterfactuals. The nomic relations holding between ‘cow’-cows and ‘cow’-horses are different in their counterfactual properties, because while ‘cow’-cows can hold without there being ‘cow’-horses relations, the reverse is not the case. If there were no ‘cow’-cows relations, there could not be nomic relations between ‘cow’ and any environmental property distinct from cows. In more simple words, any nomic relation between non-cows and ‘cow’, is parasitic on there being a ‘cow’-cows relation.

Fodor claims that his solution to the disjunction problem is non-question-begging because it is based on dependencies between nomic relations of properties, relations which are compatible with a naturalistic viewpoint and do not need to be formulated in symbolic terms. In sum, according to Fodor what makes a token a mental symbol is that it bears a nomic relation with a certain environmental property which constitutes its referent, insofar as any additional relation holding between the symbol and properties of the environment depends asymmetrically on the relation with its referent.

Taking together Fodor’s solutions to the not-all problem and the disjunction-problem, we can now sum up. According to Fodor’s view, what distinguishes us from artefacts and other entities that lack mental symbols is that we can bear

the right kind of nomic relations with certain environmental properties. Those relations are mediated by sustaining mechanisms that go far beyond perception, since they involve background knowledge and theories we have about the world. But since for the purposes of fixing the mind/world relation those mechanisms can be specified in computational terms, no appeal to other symbolic contents is required. And importantly, not any nomic relation between symbols and referent will do for the purposes of fixing content. The relation has to be asymmetrical, in the sense that any other relation between the mental symbol and environmental properties must depend, or be parasitic, on the relation between the mental symbol and its referent.

3. A CRITIQUE

In this section I put forward a critique to Fodor's notion of sustaining mechanisms which, as we have seen, constitute an integral part of his theory of content. After that, I consider some possible replies.

Sustaining mechanisms correspond to background commitments (i.e. stored theoretical knowledge) which take part in the job of fixing the nomic relations that mental symbols bear with their referents. As commented in 2.1, Fodor (1987) anticipates the objection that this could be question-begging by arguing that what fixes the nomic relation are not the mental symbols that make up the theoretical knowledge, but the causal route they sustain between the symbol and its referent. Then the structure of the theory that mediates the symbol/world correlation is supposed to be somewhat separable from the contents of its mental symbols, and to make this detachment plausible Fodor appeals to the distinction between the syntactic and semantic components of computer systems. In this case, he claims, the structure of theories responsible for establishing the symbol/world correlations that fix the symbol's contents can be specified in purely syntactic, computational terms. Thus, according to Fodor:

The picture is that there's, as it were, a computer between the sensorium and the belief box, and that the tokening of certain psychophysical concepts eventuates in the computer's running through certain calculations that in turn lead to tokenings of 'proton' (or of 'horse' or whatever) (p. 123).

These computations involve inferences drawn over true beliefs, but what fixes the belief's contents is their relation with the world, not with other beliefs involved in the inferences. A consequence of this view is that people's theories containing false beliefs (such as one stating that protons are governed by angels) would also

be capable of delivering mental symbols, insofar as they ensure a reliable symbol/world correlation (e.g. between ‘proton’ and protons). Fodor sees this as an advantage of this view since it makes it possible that different people with disparate theories about the world can share the contents of their mental symbols. This idea of distinguishing between mechanisms that enable the fixation of content and those that determine the contents themselves, is further developed on Fodor’s later work (e.g. 1998). In sum, what fixes the content of a mental symbol is the nomic relation it bears with the property of the environment it denotes, and even though the mechanisms that sustain or fix that nomic relation are required for having content, they are not relevant for determining the nature of the contents themselves.

However, I believe that it is problematic that Fodor is committed to the view that sustaining mechanisms for determining content are irrelevant in this way, because they seem to be as crucial as the nomic correlations themselves for explaining what makes genuine mental symbols possible. Indeed, in a sense, the asymmetric dependence theory just restates (from an informational perspective) something we already knew: that the relation between a symbol and its referent enjoys primacy over any other relation the symbol might have with the world. But the job of explaining why mental symbols can hold relations of this sort falls on the mechanisms that make those relations possible. Therefore, Fodor’s project of articulating “in nonsemantic and nonintentional terms, sufficient conditions for one bit of the world to be about (to express, represent, or be true about) another bit” (p. 98), needs not just to put forward the asymmetrical-dependence constraint, but also to explain what is special about the symbolic structures of mental agents, such that they can satisfy this constraint and so possess semantic properties, while other computing agents that also process information from the environment lack those properties.

To elaborate the argument, consider the following example⁶. Long before the time of Newton, mariners knew that there was a correlation between the rise and fall of the tides, and the position and phase of the moon. But a complete account of that correlation –even if the account involves nomic regularities and counterfactuals– would be insufficient to explain the tides. The mariners had no knowledge of the causal connection between the moon and tides, and to whatever extent they thought they had an explanation, it had probably to do with Gods’ benevolence or some other sort of supernatural mechanism. It was not until Newton elucidated the causal connection that we had a proper explanation of the tides.

⁶ I borrow this example from Salmon (1989: 47) who presents it as a counterexample to the deductive-nomological model of explanation. The basic idea is scientific explanations based on nomic correlations are at best incomplete without an account of the causal mechanisms underlying these correlations.

The same moral can be extended to Fodor's account of the nature of mental symbols. He offers an explanation based on (asymmetrically dependent) symbol/world nomic correlations. But as it happens with the tides, this explanation is poorly illuminating about what grounds these correlations. The sustaining mechanisms appear to be required to furnish them, but since Fodor's account quantifies over all the possible sustaining mechanisms that could fix those correlations, his explanation is left incomplete. Instead, it would be much more illuminating about the nature of mental symbols to know the computational principles or limits under which the sustaining mechanisms operate. Let me illustrate why the omission of these mechanisms is important with another example.

Imagine that in the future, scientists are able to build a robot that bears mental symbols. Even though this is still science fiction, it can be taken as a working hypothesis of the computational theory of mind that a robot endowed with the computational architecture and information processing capacities of the right complexity should be capable of thinking (see e.g. Copeland, 1993). Following Fodor's view, the scientists should at least have equipped the robot with perceptual systems and central computational mechanisms capable of reliably connecting its inner symbolic structures with their referents, in a way that fixes the appropriate nomic relations between them. So if the robot is able to think, say, about horses, it would need to possess computational mechanisms for sustaining a 'horse'-horse correlation. And in order to be genuinely semantic, that correlation would have to satisfy the asymmetrical-dependence constraint.

But as the example shows, what makes the robot capable of bearing mental symbols is not just its capacity to relate symbolic structures with its referents in the appropriate nomic way, but also that it has the right computational architecture, viz. the appropriate sustaining mechanisms. Fodor could contend that when it comes to the metaphysics of mental symbols, what matters are the semantic relations and not the sustaining mechanisms, which are just an "engineering" fact about the robot with no implications for a theory of mental symbols (Fodor, 1998: 78). However, I believe that the engineering problem is quite relevant for the purposes of drawing a line between computing agents that are capable of manipulating mental symbols and computational agents that are not. Just consider counterfactual situations: if the sustaining mechanisms had not been within certain computational limits, the robot's symbolic structure could have had a different content or no content at all. There appears to be some minimal constraints in the computational architecture of a system that are crucial for its capacity to support semantic relations.

What can we say about those minimum computational constraints? It certainly goes beyond the scope of this paper to offer a detailed account of them. However, it can be instructive to take a brief look at two types of sustaining mechanisms, understood as computationally-driven inferences, which can be found in perceptual and categorisation systems.

According to a computational approach to perception, inferential processes sustain the causal route of the information flow from the environment to the mind. These mechanisms typically consist of a series of computational processes that end up with information packed in a format suitable for tokening mental symbols. One constraint that characterises these vertical mechanisms is their sensitivity to distal properties of the environment, coded under the form of perceptual invariants. As Fodor himself has noted (Fodor 1983), “raw” information taken in through the senses is variable and mathematically insufficient to recover data about distal properties; so in order to generate the right representation of distal properties, perceptual systems have to constraint the possible interpretation of the sensory input in order to yield a constant perception of those properties. This constitutes the first stage in perception where properties of the environment are detected or individuated, and it might be a fundamental constraint for possessing a mind (Burge, 2010). Even though more than one inferential route might do the job of coding perceptual invariants, the overall process can be regarded as a perceptual sustaining mechanism that explains what makes for tokening genuine mental symbols.

When it comes to higher-order cognitive competences of categorisation, different theoretical constructs have been put forward to explain how we classify the perceptual input as falling under certain category, viz. prototype, exemplar and theory models (Laurence & Margolis, 1999). For example, exemplar models claim that when categorising, we compute the similarity between a set of exemplars and the representation of the object to be categorised. Even though prototype, exemplar and theory models have been traditionally understood as competing alternatives, it has recently been claimed these three models might coexist in human cognition as three distinct causal mechanisms of categorisation we use at different times (Machery, 2010). If this thesis is right, then these models correspond to a core set of computational mechanisms that sustain the generation of complex mental symbols.

But as remarked above, it is not the aim of this paper to go into detail about these minimum constraints on cognitive architecture. My goal in this section is ultimately to show that the general mechanisms that sustain the possession of mental symbols can be specified, and that this specification appears to be relevant for an account of what draws the line between information-processing agents with and without mental symbols. Without such an account, Fodor’s informational approach to content remains unsatisfactory as an explanation of the nature of mental symbols.

4. SOME MORE OBJECTIONS AND REPLIES

It could be objected that sustaining mechanisms do not constitute a metaphysically necessary condition for content, given it is conceivable that,

say, angels, could bear semantic relations with properties in the world without the mediation of any sustaining mechanism. But at least they surely form a nomologically necessary condition; creatures on the earth as we know it cannot bear semantic relations unless their computational architectures satisfy some minimum complexity constraints. Therefore, when faced with the question of what the minimal conditions for having mental symbols in nomologically possible creatures are, spelling out the sustaining mechanisms (represented by computational constraints) required to meet those conditions become an integral part of the response.

Another objection could be that trying to find out the sustaining mechanisms that underlie mental symbols is not really a philosophical issue, but one that belongs to empirical psychologists. After all, the sustaining mechanisms that allow us to map environmental properties onto mental symbols are just that, mechanisms, the correct description of which is an empirical enterprise rather than a conceptual task for philosophers. However, I do not believe we need to take this worry seriously, at least at the current state of cognitive science. The project of understanding what the mind is and how it works is interdisciplinary and informed by both philosophy and psychology (among other disciplines). Indeed, many philosophers studying causal theories of content –at least from Kripke– have addressed the question of the mechanisms that make mental symbols possible (cf. Jylkkä, 2008). As Dretske says in the spirit of philosophical naturalism, “if you want to know what intelligence is, or what it takes to have a thought, you need a recipe for creating intelligence or assembling a thought (or a thinker of thoughts) out of parts you already understand” (Dretske, 1994: 468-469). Without such an understanding, a theory of content is at best incomplete.

A Fodorian might want to press this point and insist that asking for details about sustaining mechanisms is otiose, given that we can always say more about mechanisms, even by appealing to different explanatory levels such as neurobiology or physics. But of course not any implementational detail is relevant for explaining a given phenomenon; explanations are always framed within certain theoretical background which demarcates which details are relevant or not. Fodor's theory of content is explicitly formulated within the paradigm of the computational theory of mind, and so to determine the constraints under which the mental symbols work is primarily a job for cognitive psychologists. And it is interesting to point out that the importance of these constraints for understanding what is distinctive of mental creatures seems to be a fundamental assumption of the computational theory of mind. Since if the mind is a Turing machine, and given that not any implementation of a Turing machine is capable of instantiating a mind, it follows that what makes for a mind is implementing a Turing machine of the appropriate complexity (see Kim, 2006: 133).

5. CONCLUSIONS

There is good reason to believe that thinking is
computation and that content is information [...]
(Fodor, 1994: 26).

Fodor's commitment to informational approaches to content has two main motivations. One is that it makes possible a naturalistic formulation of how mental tokens could have intentional content. In other words, it can explain how content is determined by appeal to notions such as nomic relations, causes and computational processes, which are compatible with physicalism and do not involve semantic properties. The second is that his informational account is compatible with semantic atomism, the view that a symbol's content is determined by its relations with the environmental property it denotes, and not by its relation to other mental symbols⁷.

After looking at what motivates Fodor's version of the informational approach, it shall be useful to return to the main tenets of his view reviewed in this paper. The asymmetric dependence theory deals with the disjunction problem by positing that the nomic link a mental symbol holds with its referent is prior to any other relations bearing between the symbol and other environmental properties. The theory implies that no symbol-symbol relationships are essential for determining content, thus preserving the atomic nature of mental symbols. The notion of sustaining mechanism, in turn, attempts to sort out the not-all problem by proposing that the causal route linking a symbol with its referent is fundamentally computational, whereas other symbols that take part in this route are relevant for determining content only with respect to their formal – computational – properties. Again, since no appeal to other symbol's semantic properties is required, the solution to the not-all problem is compatible with naturalism and semantic atomism.

I believe that Fodor's informational account is consistent with his mentioned motivations, and, in this paper, I do not want to challenge semantic atomism or its naturalistic commitments. Indeed, I regard his refinements to informational approaches as illuminating about the sort of symbol-world nomic relation that might ground genuine semantic properties in a way that is compatible with atomism. However, I contend that his view is at best incomplete about the computational and information processing requirements for instantiating symbolic structures with the

⁷ Fodor's reasons for endorsing atomism are mostly related to his fear of semantic holism and its undesirable consequence of the (alleged) indeterminacy of content. See Fodor and Lepore (1993) for discussion.

relevant nomic relations. This renders Fodor's view unsatisfactory when it comes to distinguishing us from other computing agents that also process information from the environment but to which the ascription of mentality would be implausible (see section 1). We could grant Fodor his claim that the semantic dimension of the theories underlying sustaining mechanisms is irrelevant for content⁸. However, a notion of sustaining mechanism that quantifies over the computational structures of the mediating theories that make the emergence of content possible, fails to shed light on what an informational approach is purported to explain. To use Fodor's expression, sustaining mechanisms might be just an engineering problem, but it is one that has to be taken seriously if we want to understand what makes us the only known computing machine capable of bearing mental symbols so far.

REFERENCES

- Adams, Frederick. "The informational turn in philosophy", *Minds and Machines*, 13 (2003): 471-501.
- Aguilera, Bernardo. "Behavioural explanation in the realm of non-mental computing agents", *Minds and Machines*, 25/1 (2015): 37-56.
- Burge, Tyler. *Origins of Objectivity*. New York: Oxford University Press, 2010.
- Copeland, Jack. *Artificial Intelligence: A Philosophical Introduction*. Oxford: Blackwell, 1993.
- Dretske, Frederick. *Knowledge and The Flow of Information*. Cambridge, MA: MIT Press, 1981.
- Dretske, Frederick. "If you can't make one, you don't know how it works." *Midwest Studies in Philosophy*, 19/1 (1994): 468-482.
- Fodor, Jerry. *The Modularity of Mind: An Essay in Faculty Psychology*. Cambridge, MA: MIT Press, 1983.
- . *Psychosemantics: The Problem of Meaning in The Philosophy of Mind*. London: Bradford, 1987.

⁸ For a defence to the view that the semantic dimension of sustaining mechanisms is constitutive of content see Rupert (2000) and Jylkkä (2009).

- . *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press, 1990.
- . *The Elm and The Expert*. Cambridge, MA: MIT Press, 1994.
- . *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press, 1998.
- . LOT 2: The language of thought revisited. Oxford: Clarendon Press, 2008.
- Fodor, J., & Lepore, E. *Holism: A Shopper's Guide*. Oxford: Blackwell, 1993.
- Jylkkä, J. W. "Theories of natural kind term reference and empirical psychology", *Philosophical Studies*, 139/2 (2008): 153-169.
- Jylkkä, J. W. "Why Fodor's Theory of Concepts Fails", *Minds and Machines*, 19/1 (2009): 25-46.
- Kim, Jaegwon. *Philosophy of Mind*. Cambridge, MA: Westview, 2006.
- Laurence, S., & Margolis, E. "Concepts and Cognitive Science". *Concepts: Core Readings*. In Eric Margolis & Stephen, Laurence (Eds.), Cambridge, MA: MIT Press, (1999): 3-81.
- Machery, Edouard. "Précis of doing without concepts", *The Behavioral and Brain Sciences*, 33/2-3 (2010): 195-206.
- Margolis, Eric. "How to Acquire a Concept", *Mind and Language*, 13/3 (1998): 347-369.
- Rupert, Robert. "Dispositions indisposed: semantic atomism and Fodor's theory of content", *Pacific Philosophical Quarterly*, 81/3 (2000): 325-348.
- Salmon, Wensley. "Four decades of scientific explanation", *Scientific Explanation*. In Phillip Kitcher & C. Wesley Salmon (Eds.), Minneapolis: University of Minnesota Press, (1989): 3-219.
- Stalnaker, Robert. *Inquiry*. Cambridge, MA: MIT Press, 1984.