

NOT ANOTHER ESSAY ON MODULARITY OF MIND:
ON THE POSSIBILITY OF MODULAR CENTRAL SYSTEMS

Sheldon Chow
schow53@uwo.ca

The University of Western Ontario
Department of Philosophy
Talbot College
London, Ontario
Canada, N6A 3K7

Area: Philosophy of Mind
Word count: 3762 (including notes)

Abstract

According to the massive modularity thesis, the so-called central cognitive capacities such as reasoning and belief formation, which are typically thought of as paradigmatic examples of unencapsulated and nonmodular cognition, are held to be subserved by a number of functionally distinct modules. But there is little consensus on what modularity amounts to. In this paper I raise the issue of whether modules are *informationally encapsulated* with an aim to show that central systems modularity ceases to be meaningful or interesting if modules aren't conceived to be so.

Introduction

There is a stream in cognitive science which holds that the human mind is *massively modular*—that for each cognitive capacity there is a distinct cognitive mechanism (usually characterized as domain-specific and informationally encapsulated) dedicated to its function. According to the massive modularity thesis, even the so-called central cognitive capacities such as reasoning and belief formation, which are typically thought of as paradigmatic examples of unencapsulated and nonmodular cognition, are subserved by a number of functionally distinct modules. However, claims about massive modularity notwithstanding, in the literature on modularity of mind one often comes across the acknowledgement that there are various differing characterizations of modules (see e.g., Carruthers, 2006a; Samuels, 2000; Segal, 1996). As a result, there is little consensus on what modularity of mind amounts to. This leaves the massive modularity thesis in a vulnerable position, since any unclarity in a theory of modularity will be shared by a theory of massive modularity. One issue on which there isn't widespread agreement is whether modules are *informationally encapsulated*.

In this paper I address the question of whether we are left with any meaningful or interesting notion of massive modularity if informational encapsulation is not considered an essential to modularity. I begin by explaining how the computational tractability argument, which is commonly invoked to claim that modules are informationally encapsulated, doesn't actually entail encapsulation. Instead, the use of heuristics can achieve the same computational frugality as what is expected of encapsulation. I next consider an account of

modular central systems suggested by Sperber (1994), and argue that the modules subserving its activities must be *uninterestingly* encapsulated at best. In doing so, I motivate a proposal to characterize encapsulation in terms of restrictions on access to kinds of information held or generated in the mind, rather than in terms of restrictions on access to certain systems in the mind. Lastly, I briefly discuss how an interesting view of modularity will preclude modular central systems, even if we insist on computational frugality rather than encapsulation as a feature of modules.

Modularity and informational encapsulation

Fodor (1983, 2000) famously uses an argument from computational tractability to claim that mental modular processes require constraints on the amount of information available for computation,¹ and hence that modules must be informationally encapsulated. The idea is that information available to a given module is restricted to that which is contained in its proprietary database—it simply can't access other systems in the mind, and therefore, if we consider the module-plus-database as composing a system, a system is encapsulated with respect to any information not contained in it. Modules can thereby compute quickly and efficiently since there is no need for information searches—no need to seek out information relevant to its computations.

Fodor is often cited as claiming that encapsulated cognitive mechanisms have access to less than all of the information possessed by an agent in the course of its computations (see for example Carruthers, 2006a). This contrasts with the kind of unencapsulation Fodor (2000) seems to have in mind for central reasoning systems with his argument from abduction—global considerations require the accessibility of the totality of one's beliefs. But access to less than

¹It is certainly by no means settled whether the mind is actually realized by computations, as certain connectionist theorists will be eager to point out. Those who do not buy into the computational theory of mind will find the argument from computational tractability unconvincing. However, it is beyond the scope of the present paper to wade in on the debate. I propose to simply run with the premise that the mind is realized by computations.

all of the information possessed by an agent isn't the entire story for what must be meant by encapsulation. For on the suggested reading every cognitive system would be encapsulated, since presumably no system has access to *all* information possessed by an agent. Rather, Fodor's full story is actually this: "to the extent that . . . systems are informationally encapsulated, of all the information that might *in principle* bear upon a problem of . . . analysis only a portion (perhaps only a quite small and stereotyped portion) is actually admitted for consideration" (Fodor, 1983, p. 70). This is to say that a cognitive system is encapsulated to the extent that it has access to less than all information that might *in principle* bear on its operations. And, of course, restrictions on a module's access to information is ensured through suitable restrictions on its access to other systems in the mind.

To be clear, information that in principle bears on the operations of a system is information that is *relevant* to the system's tasks. To illustrate, consider the most plausible case for encapsulation: the low-level visual processes, such as colour perception or length perception. The paradigmatic example is the Müller-Lyer illusion which seems to indicate that the mechanisms responsible for visual length perception do not have access to the belief that the two lines are in fact the same length. But this belief is relevant to visual length perception. Hence, such information in principle bears on the problems of analysis of visual processing, but is not contained in the proprietary database upon which the modules responsible for visual processing draw. And it is because there exists information somewhere else in the mind that can in principle bear on the analysis tasks of visual processing, but to which it seems it has no access, that we think that the system is encapsulated. We don't believe the system to be encapsulated merely because it has access to less than all information held in the mind. For even the central reasoning system doesn't have access to all information held in the mind. For instance, central reasoning seems to have no access to the preconceptual information of low-level language processing. But being encapsulated with respect to this preconceptual

information simply doesn't matter for what the central reasoning system is supposed to do. What matters is that the central reasoning processor have access to all the *beliefs* and other *conceptual* information within an agent's mind. These latter kinds of information are those which are relevant to, and so can in principle bear on central reasoning. Preconceptual low-level language information, on the other hand, is generally irrelevant to its processes; it's simply not the right *kind* of information with which central reasoning is supposed to work. And it is because it has access to all the right kinds of information that the central reasoning system is often claimed to be unencapsulated, not because we believe it to have access to *all* information held in the mind.

Nevertheless, Samuels (2005) claims, *pace* Fodor, that the computational tractability argument doesn't entail informational encapsulation. Instead, Samuels asserts that all that computational tractability requires is that computations be suitably *frugal*, and that appropriately frugal computational processes can be realized through heuristics.² Heuristics instantiate search rules and stopping rules which limit the amount of information that actually gets surveyed in computations, thereby avoiding exhaustive searches.³ What is implied for modular processes is that it is possible for any information held within the system that is relevant to a module's computational processing to be brought to bear on it; it is just that, through the use of heuristics, most of the information merely *won't* bear on the module's computations at a given time.⁴ In other words, a module may be unencapsulated but yet

²A similar view can be found in Carruthers (2006a). But Carruthers actually believes that we may still understand modules to be encapsulated in the sense of most of the information held in the mind can't bear on the system in question in the course of its processing. However, Samuels later criticizes Carruthers' view, claiming that "not only is [Carruthers' brand of encapsulation] different from what most theorists mean by 'encapsulation,' but it's simply what you get by denying exhaustive search; and since virtually no one thinks exhaustive search is characteristic of human cognition, the present kind of 'encapsulation' is neither distinctive nor interesting" (Samuels, 2006, p. 45).

³The most notable researchers on heuristics commonly employed by the human mind is Gigerenzer and his colleagues (Gigerenzer, Todd, & the ABC Research Group, 1999).

⁴Computational frugality also means, in addition to limiting the amount of information available for computation, instantiating algorithms that are not intractably complex. Furthermore, tractability would also require that a system maintain sufficient processing capacity to carry out its computations. Though these are important aspects of computational frugality and tractability, I will leave them to one side, since my main focus and concern is for limitations

computationally perform *as if* it was encapsulated. The upshot, then, is that the computational tractability argument isn't sufficient for holding that informational encapsulation is a necessary feature of modules.

None of this of course shows that the tractability argument entails that no system may be encapsulated. Indeed, encapsulated systems generally are computationally frugal by their very nature. Thus, the frugality constraint just requires that the amount of information bearing on mental computations is in some way limited, whether achieved through encapsulation or through heuristical processes. It should not be expected, therefore, that no system of the mind will be encapsulated. And, in fact, there is considerable empirical evidence that there are a number of encapsulated mental systems.⁵ Samuels (2006) emphasizes, however, that the systems for which we have the most empirical evidence to suggest the plausibility of encapsulation are the peripheral, non-central systems of cognition (see also Samuels, 1998, 2000; Fodor, 1983, 2000). And indeed, it seems as if the farther we enter into the conceptual domain, the less it seems plausible that encapsulated mechanisms underlie our mental processes.

Nevertheless, many continue to argue that all or most of our cognition is modular in character—that the mind is *massively modular*—implying not only that our peripheral systems are modular, but so are our *central* systems (paradigmatically those responsible for our reasoning and decision making).⁶ Central systems are at the focus of the debate between massive modularists and anti-massive modularists because such systems just don't seem to be

in the amount of information bearing on computations.

⁵In addition to the low-level visual processes, which was said above to present the most plausible case for an encapsulated system, Fodor (1983) suggests that low-level language processes, such as phonological processing, make for a good case for informational encapsulation. Samuels (2006) claims that there are other plausible cases indicating encapsulation (though not as plausible as the low-level visual or language processes), such as face recognition and approximate arithmetic.

⁶For example, it has been suggested that there are modular mechanisms for such central processes as probabilistic inference (Gigerenzer, 1994, 1996), biological categorization (Pinker, 1994, 1997), social reasoning (Cosmides & Tooby, 1992), and the so-called “theory of mind” inference (Baron-Cohen, 2005).

candidates for modularity. Central systems admittedly allow for free exchange of vast amounts of information, any and all of which can in principle bear on their operations.⁷ Dreams of snakes, for instance, can bear on theorizing about molecular structures of chemicals. Massive modularists who recognize the informational-encompassing character of central systems may respond in a variety of ways. Carruthers (2006a), for instance, attempts to redefine encapsulation, claiming that we can retain a “wide” sense of encapsulation through heuristics, since most of the information held in the mind is prevented from bearing on a system in the course of its processes. However, the resulting concept isn’t really a encapsulation at all.⁸ Alternatively, massive modularists might concede that *some* systems of the mind are likely not modular—a position that Samuels (1998, 2000) calls “weak massive modularity”.⁹ Yet, “weak massive modularity” seems to be a misnomer since it really takes the “massive” out of the massive modularity thesis. Some massive modularists, however, attempt to cast the workings of central systems in terms of the collective effect(s) of the operations of a suite of encapsulated modules. This last option is taken up by Sperber (1994).¹⁰

Against central systems modularity

The story Sperber tells as an explanation of how central processes can be subserved by modules begins with the assertion that the apparent free exchange of information implies only that one particular modular architecture for central systems can’t be right, viz., one which is constituted by a number of mutually unconnected modules, where information possessed or computed by one won’t “find its way” to the others. “If, on the other hand,” as Sperber continues, “the output of one conceptual module can serve as input to another one, modules can each be informationally encapsulated while chains of inference can take a conceptual

⁷This is a feature acknowledged by those who deny central systems modularity (e.g., Samuels, 2000, 2006) as well as those who advocate for it (e.g., Carruthers, 2006a; Sperber, 1994).

⁸See footnote 2 above.

⁹This is the position that Samuels (2000, 1998) suggests Cosmides and Tooby (1994) actually take.

¹⁰See also Pinker (1997).

premise from one module to the next and therefore integrate the contribution of each in some final conclusion” (Sperber, 1994, p. 48). Though coherent, this story of central systems modularity is ambiguous. Specifically, it is not clear whether the input information received by a given module can be the output of (i) only a single other module, (ii) a select few modules, (iii) many (not all, but more than a few) modules, or (iv) all other modules in the system.

Option (iv) is likely not the story Sperber wants to tell, since it seems to describe a suite of *unencapsulated* modules. Further, while it is possible for a system to be set up and operate as (i) or (ii) describes, it is unlikely that either is how the central systems of humans are actually set up and operate. First, the sheer speed of our thinking and reasoning requires that the connections between mechanisms be set up to operate in parallel rather than serial fashion. What this means is that more connections are needed between numbers of mechanisms than simply a one-to-one or few-to-one output-input interface to facilitate information flow. Second, even if this speed issue can get resolved (though heuristics, maybe), as often stressed one of the most crucial and interesting features of central processes is the *free flow of information*, where *any* relevant information can bear on its operations. What is described in (i) and (ii) doesn’t seem to provide room for this feature. For instance, if the mechanisms subserving central reasoning processes are not set up in the right way, certain conceptual information may not be permitted to be brought to bear on their operations; if not set up just so, dreams of snakes may not be able to be brought to bear on theorizing about molecular structures of chemicals. But such cannot be possibilities of central systems. What seems to be required, then, is a sufficient amount of certain kinds of connections between the mechanisms which subserve central systems. The connections of (i) and (ii) appear to be insufficient.

On the other hand, (iii) can accommodate the way central systems actually operate. According to this interpretation, information would get integrated as it flows through the modules comprising the system, and it would all come together somewhere to reach what

Sperber calls the “final conclusion”. This final conclusion is presumably just the final output of the central system. This is probably the best case that can be made in favour of Sperber’s story of modular central systems. But does this deliver what it is supposed to deliver? To see that the answer to this question is an equivocal, though warranted, “probably not”, let us take a closer look at what informational encapsulation really amount to.

As indicated above, informational encapsulation is usually characterized in terms of restrictions on access to other systems in the mind. However, at bottom, informational encapsulation is really about restrictions on access to certain kinds of information. Imagine for instance two systems A and B.¹¹ System A conducts searches across all stored beliefs in the course of its operations, and hence isn’t encapsulated (or at least isn’t interestingly encapsulated). System B, on the other hand, doesn’t conduct searches across all stored beliefs, but queries a range of other systems, which collectively search all stored beliefs on its behalf. Certainly it would be absurd to say that while System A isn’t encapsulated, System B really is encapsulated since all of the information made available to it comes from a limited set of systems. Encapsulation doesn’t only concern access to other systems, but access to kinds of information. And it is access to information, as this example illustrates, that’s at the heart of encapsulation. For notwithstanding that System B is isolated from certain other systems, it essentially has access to all information that might in principle bear on its operations since it is not in any meaningful sense isolated from information held or generated elsewhere in the mind; hence, it is not in any meaningful sense encapsulated.

The moral for Sperber’s story of modular central systems can be made clear if we take a close look at the flow of information in his system. In order to facilitate free flow of information, it must be the case that each mechanism in the system is properly connected to

¹¹This example is an adaptation of one given by Samuels (2005, p. 113), and it is very similar to a problem Carruthers (2006a, p. 6) acknowledged in a different context. I believe that the moral remains the same for both authors, but it’s interesting to note that Carruthers draws different conclusions from it than what Samuels does. Ironically, the moral militates against Carruthers’ own massive modularity inclinations.

each other in such a way that enables the appropriate kinds of information to be integrated, or otherwise be brought to bear on the system's processes. But information flow has to do with more than just connections between cognitive mechanisms. It has to do with how information gets shipped around between mechanisms—what outputs can and do serve as input for which mechanisms. Consider, for example, central reasoning where it is thought that any and all beliefs contained in the mind can in principle bear on its processes. It simply doesn't matter if all such information is received from only a limited set of modules. Whether its operations occur within a single mechanism or through a suite of mechanisms, the information that may bear on central reasoning can potentially be any set of beliefs whatsoever. Thus, the central reasoning processor must in fact *have* access to all the kinds of information which can possibly bear on its operations, and so the appropriate mechanisms subserving its activities can't be in any meaningful sense isolated from any belief held or generated in the mind. Thus, Sperber can tell his central systems story where certain modules will be able to receive as input the outputs of one or more (but not all) other modules. But at least some modules will have to be at best *uninterestingly* encapsulated, since central systems require the free flow of information, which means that the appropriate underlying mechanisms cannot be isolated from any of the kinds of information that can in principle bear on their operations.

Modularity worth wanting

Free flow of information is precisely what provokes the idea that central systems are unencapsulated. As Samuels (2005) points out, encapsulation is *architecturally imposed*. What he means by this is that encapsulation is a special sort of restriction on information access. It is special because encapsulation is a relatively enduring characteristic of the system, it is not a mere product of performance features of the system (e.g., lack of computational resources), and it is a property of the system that is "cognitively impenetrable" (Pylyshyn, 1984) (roughly, it

cannot be changed by psychological processes alone, such as mere alterations of beliefs and desires) (Samuels, 2005, p. 112). An implication of this, I take it, is that the extent to which a mechanism is encapsulated depends on the way it is hardwired within the network of the mind as a whole—a system is encapsulated in virtue of it not being connected in certain ways to certain other systems in the mind, preventing the exchange of certain kinds of information. And this is just to say that encapsulation is a product of restrictions on the flow of information within a system. If the architecture of a system enables a great enough flow of information, the system may become *conceptually indistinguishable* from an unencapsulated mechanism, having access to all (or nearly all) information which can in principle bear on its operations.

But can heuristics save central systems modularity? If we require only that modules be frugal in their computations, we can envision a suite of modules subserving central systems, at least some of which having complete access to any and all information which might in principle bear on its operations, but recruiting the appropriate heuristics to extract and integrate only a certain amount of certain kinds of information through its processing. None of what has been said about the free flow of information of central systems precludes this possibility.

Nevertheless, computational frugality through heuristics is not enough to save central systems modularity. The reason for this is that, regardless of how computational frugality is achieved, it's informational encapsulation that makes modularity interesting.¹²

It is true that the system can be described as being composed of a number of distinct and isolable parts, and these parts may in some sense be called modules, but this view is by no means controversial (Samuels, 2006). An interesting theory of cognition must go beyond characterizing modules simply as functionally distinct, isolable parts or components of the mind.¹³ In particular, an interesting theory of modularity must have something interesting to

¹²Cf. Fodor: “it's informational encapsulation, however achieved, that's at the heart of [the modularity thesis]” (Fodor, 2000, p. 63).

¹³See Carruthers (2006a, pp. 1-7) for an opposing view.

say about the flow of information within the mind.¹⁴ Thus, a central issue that must be addressed is the extent to which cognitive mechanisms are able to exchange information with others. This means that an interesting theory of cognition must address the extent to which cognitive mechanisms are informationally encapsulated. And since informational encapsulation admits of degrees (Carruthers, 2006a; Samuels, 2005), what is of concern is whether we possess cognitive mechanisms that are encapsulated (again) to some *interesting* extent (Samuels, 2005). For, as we saw, a system may be encapsulated alright, but uninterestingly so.

Conclusion

I don't believe anyone will deny that at least some of human cognition is modular, especially that of the peripheral systems. But the case isn't so clear when it comes to central cognition. What I've shown is that an interesting conception of modularity requires some interesting notion of encapsulation, but such a conception of modularity precludes the way we understand central systems to operate. Central systems require free flow of information, which means that they must be connected in the right sort of ways with other systems of the mind and within the architecture of the mind as a whole. As a result, a modular view of central systems is at best uninterestingly encapsulated, and so uninterestingly modular. This is precisely why central systems are the usual suspects for nonmodular cognition. Even if our view of modularity jettisons encapsulation as a core feature of modules, but maintains minimally that their computations be sufficiently frugal, we won't be able to save a modular view of central systems. For informational encapsulation is what makes modularity interesting. Thus, we aren't left with an interesting view of massive modularity.

References

Baron-Cohen, S. (2005). *Mindblindness*. Cambridge: The MIT Press.

¹⁴This is recognized not only Samuels (2000), who believes central systems to be nonmodular, but also Sperber (1994), who believes that central systems can be characterized as modular.

- Carruthers, P. (2006a). *The architecture of the mind: Massive modularity and the flexibility of thought*. Oxford: Oxford University Press.
- Carruthers, P. (2006b). The case for massively modular models of mind. In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (p. 3-21). Malden MA: Blackwell Publishing.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (p. 163-228). New York: Cambridge University Press.
- Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: The evolution of functional organization. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (p. 85-116). Cambridge: Cambridge University Press.
- Fodor, J. (1983). *The modularity of mind*. Cambridge: The MIT Press.
- Fodor, J. (2000). *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge: The MIT Press.
- Gigerenzer, G. (1994). Why the distinction between single-event probabilities and frequencies is important for psychology (and vice versa). In G. Wright & P. Ayton (Eds.), *Subjective probability* (p. 129-161). New York: John Wiley.
- Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review*, 103, 592-596.
- Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.
- Gigerenzer, G. (2001). The adaptive toolbox. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (p. 37-50). Cambridge: The MIT Press.
- Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (p. 62-88). Malden, MA.: Blackwell Publishing.
- Gigerenzer, G. (2006). Bounded and rational. In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (p. 115-133). Malden, MA.: Blackwell Publishing.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group (Eds.). (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Pinker, S. (1994). *The language instinct*. London: Penguin Books.
- Pinker, S. (1997). *How the mind works*. New York: W. W. Norton.
- Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge: The MIT Press.
- Samuels, R. (1998). Evolutionary psychology and the massive modularity hypothesis. *The British Society for the Philosophy of Science*, 49, 575-602.
- Samuels, R. (2000). Massive modular minds: Evolutionary psychology and cognitive architectures. In P. Carruthers & A. Chamberlain (Eds.), *Evolution and the human mind* (p. 13-46). Cambridge: Cambridge University Press.
- Samuels, R. (2005). The complexity of cognition: Tractability arguments for massive

- modularity. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and contents* (p. 107-121). Oxford: Oxford University Press.
- Samuels, R. (2006). Is the human mind massively modular? In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (p. 37-56). Malden MA: Blackwell Publishing.
- Segal, G. (1996). The modularity of theory of mind. In P. Carruthers & P. K. Smith (Eds.), *Theories of theories of mind* (p. 141-157). Cambridge: Cambridge University Press.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (p. 39-67). Cambridge: Cambridge University Press.