Modular architectures and informational encapsulation: A dilemma

Dustin Stokes (University of Toronto)
Vincent Bergeron (University of Ottawa)

Abstract

Amongst philosophers and cognitive scientists, modularity remains a popular choice for an architecture of the human mind, primarily because of the explanatory value of this approach. Modular architectures can vary both with respect to the strength of the notion of modularity and the scope of the modularity of mind. We propose a dilemma for modular architectures, no matter how these architectures vary along these two dimensions. First, if a modular architecture commits to the *informational encapsulation* of modules, as it is the case for modular theories of perception, then modules are on this account impenetrable. However, we argue that there are genuine cases of the cognitive penetrability of perception and that these cases challenge any strong, encapsulated modular architecture of perception. Second, many recent massive modularity theories weaken the strength of the notion of module, while broadening the scope of modularity. These theories do not require any robust informational encapsulation, and thus avoid the incompatibility with cognitive penetrability. However, the weakened commitment to informational encapsulation significantly weakens the explanatory force of the theory and, ultimately, is conceptually at odds with the core of modularity. We conclude by proposing a non-modular, but explanatorily sufficient, notion of *functionally independent system*.

Amongst philosophers and cognitive scientists, modularity remains a popular choice for an architecture of the human mind. Jerry Fodor (1983), who was influential in establishing the concept of a module, writes: "One day . . . Merrill Garrett made what seems to me to be the deepest remark that I have yet heard about the psychological mechanisms that mediate the perception of speech. 'What you have to remember about parsing is that basically it's a reflex.'" (Dedication). Reflexes are quick, inflexible, involuntary responses to stimuli, and Fodorian modules are reflexes. In its most recent and general form, the modularity hypothesis consists in viewing the human mind, or at least part of it, as a configuration of quick specialized mental mechanisms, or subsystems, that are functionally independent of one another, and that typically operate over a distinct domain of information.

There are compelling theoretical and empirical motivations for this approach. Theoretically, modularity nicely accommodates adaptationist and other evolutionary explanations of mental phenomena. It also provides materials for a simple explanation of important empirical data, including a wide range of behavioural dissociations, as well as the speed and robustness of processing enjoyed by the human mind. Most broadly, modularity provides an intuitive framework for characterizing the structure-function correspondence between brain structures and particular perceptual and cognitive functions.

Although it is sometimes misrepresented as doing precisely this, Fodor's pioneering discussion of the concept did not involve a definition of 'module'. (Fodor 1983; see also Coltheart 1999). Fodor did, however, provide a list of properties symptomatic of modules. Fodorian modules are *typically* domain specific, hardwired, computationally autonomous, informationally encapsulated, fast, and their operation is mandatory. It is noteworthy how much of this characterization follows the *reflex* metaphor. Domain-specificity parallels the singularity of the stimulus that sets off a reflex; autonomy, mandatoriness, hardwiring, and encapsulation mirror the standard reflex-arc model. Fodor maintains that "The notion of modularity ought to admit of degrees" (Fodor 1983: 37), and that "if a psychological system has most of the modularity properties, then it is very likely to have all of them" (Fodor 1983: 137).

Importantly, Fodor claimed only that *input systems* are modular. His primary subject matter was perceptual systems, but he also made the case for systems devoted to low-level linguistic decoding. Higher-level conceptual or cognitive systems, then, are not modular on Fodor's general architecture.

Commitments with respect to Fodor's original analysis of modularity vary. Several modularity theorists take domain-specificity to be definitive of modularity (Coltheart 1999). Fewer require innate specificity, even if related explanations and arguments often invoke evolutionary considerations. Others maintain that modules are informationally encapsulated and computationally autonomous (Farah 1994, Sperber 1996; 2001.)[1] Recent theorists have extended the modularity thesis beyond Fodor's input systems. It is common among evolutionary psychologists to endorse some version of what Dan Sperber (1994) has called the *massive modularity thesis*. The general hypothesis states that all, or nearly all, of the mind is modular, and modules have been postulated to account for cognitive capacities as diverse as theory of mind, face recognition, cheating detection, reading, and a variety of social understanding abilities.

Our suggestion is that, in spite of the varied commitments in the modularity literature, *informational encapsulation* is essential to a distinctive, non-trivial modularity theory. As it will be understood here, if a module *m* is informationally encapsulated then *m* cannot, during the course of its processing, access or compute over information found in other components of the overall system. As such, an

[1] In at least two places, Fodor himself explicitly states that "informational encapsulation is an essential property of modular systems" (Fodor 1985: 3; see also 1983: 71). Though elsewhere he is less clear on his commitment regarding the same claim.

encapsulated module *m* is *impenetrable* with respect to the other components of the system, since the processing of *m* is insensitive to (and so does not compute over) the information available elsewhere in the system (Pylyshyn 1980). This basic analysis of modularity is important to any substantive modular account of the mind because it constitutes the foundation of modularity in so far as modules are, in a sense to be explained later, *functionally independent systems*.

In this respect, the modularity theorist faces a dilemma that hinges on the commitment to informational encapsulation. On the one hand, a commitment to informational encapsulation, as made by modular theories of perception, is inconsistent with the cognitive penetration of perceptual experience. And, we argue, there are genuine cases (old and new) of the cognitive penetrability of perception. On the other hand, as recent modularity theorists have done, one might weaken the notion of 'module' so as not to require any robust informational encapsulation. The result, however, is an account that is inconsistent with one of the central motivations for modular architectures and, more fundamentally, with the conceptual core of the very notion of modularity.

The first horn challenges strong, *encapsulated modularity*: any modularity theory that includes a commitment to informational encapsulation. The second horn challenges *massive modularity*, which broadens the scope of modularity but weakens the notion of modularity so as not to require informational encapsulation. Either way, modular architectures appear to be significantly challenged.

We conclude by proposing a non-modular, but explanatorily sufficient, notion of functionally independent systems. This notion characterizes the basic architectural units of the mind in a way that both avoids the challenges to modular architectures and, nonetheless, explains the very phenomena that motivate modular architectures.

I. Informationally encapsulated modules: Cognitive penetrability and the challenge for encapsulated modularity

Both encapsulated and unencapsulated modularity theorists take perceptual systems to be modular. If perceptual modules are informationally encapsulated, then at the very least, they are not penetrable by the information or processing of higher-level cognitive systems. Most theorists seem to take the concepts 'informational encapsulation' and 'cognitive impenetrability' to be co-extensive, if not equivalent—Fodor in fact originally argued for the encapsulation of modules by arguing against claims about the cognitive penetrability of input systems (Fodor 1983: 73-86). The following discussion requires only the assumption that informational encapsulation *of perceptual modules* entails cognitive impenetrability.

On this account, then, perceptual processing is not influenced by cognitive states like beliefs or desires. Evidence of cognitive penetration of perception thus threatens any modular theory that includes a commitment to informational encapsulation of perceptual modules.

A cognitive impenetrability thesis, regarding perceptual experience, says that for any two perceivers, if one holds fixed the object or event of perception, the perceptual conditions, the spatial attention of the subject, and the conditions of the sensory organ(s), then the perceptual experiences of those perceivers will be identical (see Macpherson, forthcoming). If the experiences of the two perceivers are distinct in these circumstances, and as a result of distinct cognitive states of the perceivers, then experience is, instead, cognitively *penetrable*.[2]

It will be useful here to offer some clarifications. First, distinguish perceptual experience from higher-level cognitive and affective states and processes like belief, judgement, desire, emotion, and so on. The former is, whatever else one says about it, characterized by phenomenal character or content and depends non-trivially on one or more sensory organ. Debates abound regarding how to draw the line between perception and cognition. The only point that need be granted here is that there are clear cases of perceptual states and clear cases of cognitive states. So there *are* visual experiences, auditory experiences, olfactory experiences, and so on; and these can be distinguished from states like belief and processes like decision making.

Second, distinguish the cognitive penetration of perceptual experience from the cognitive penetration of perceptual processing. The former concerns some difference in the phenomenal content or character of a perceptual experience, where this difference depends non-trivially upon some cognitive state or processing in the system. The latter only concerns some cognitive effect on perception at the level of processing. The fact that processing at *some* stage is cognitively penetrated does not, by itself, entail the cognitive penetration of experience. Experience may depend on a wider class of processing and, in principle, the cognitive influences on perceptual processing (at some particular stage or other) may not ultimately influence conscious

[2] For an alternative definition, see Siegel, forthcoming A. For a definition of one particular type of cognitive penetration—desire-influenced perception—and how it engages with the relevant debates concerning cognitive penetrability and theory-ladenness, see [AUTHOR MASK B-XXXX].

experience. Moreover, perceptual processing may not give rise to a conscious experience but rather, for example, the sub-personal guidance of motor performance (Goodale and Milner 1992).

However, cognitive penetration of perceptual experience does entail cognitive penetration of perceptual processing *at some level*. There is much to be said here. The only assumption we need regarding the relation between perceptual experience and perceptual processing is this. Whether one takes experience to be identified with, constituted by, supervenient upon, or the output of perceptual processes, a difference in perceptual experience implies a difference in perceptual process. This will be generally true—albeit for different reasons—no matter how one's metaphysics of mind varies according to these alternatives. So if experience is penetrated, then information relevant to cognitive systems, or the processing of cognitive systems, directly influences the processing of perceptual systems. It is this entailment relation that is important for the criticism we offer below.

One final point: Although the cognitive penetration of experience entails the penetration of processing at some stage, the cognitive penetration of experience is compatible with the cognitive impenetrability of processing at some stage or even most (but not all) stages. This point is instructive: one cannot argue from the purported fact that some particular perceptual module is impenetrable to the claim that perception broadly or perceptual experience (in that modality) is impenetrable. Zenon Pylyshyn, for instance, argues that "early vision" is not penetrable by cognition (Pylyshyn 1999). Pylyshyn's empirical claim about early visual processing, even if true, is insufficient to support the thesis that perception is cognitively impenetrable. Indeed, Pylyshyn admits that the output of this component of the visual system, as he and most theorists understand it, does not (alone) determine perceptual experience. His defence of cognitive impenetrability is thus consistent with the cognitive penetrability of perception; one might accept that the computations performed by the early visual system are impenetrable by cognitive states but maintain that perceptual processing is penetrated elsewhere such that the resulting perceptual experience is causally dependent upon cognition[3]

In the middle 20th century, the New Look movement in psychology, in a slogan, theorized perception as an active construction of representations of the environment. Importantly, perceptual representations are constructed in a way informed by the perceiver's "mental set"—her expectations, needs, values, desires, and other higher-level states. Although initially very influential, this New Look model of perception has been largely dismissed. As a result, plausible experimental evidence for the cognitive penetrability of experience, as provided by New Look studies, was also dismissed. This dismissal is unwarranted. Some of the classic research done by these psychologists evades the standard strategies employed to reject it, and in turn remains some of the best existing evidence for cognitive penetrability of experience. (And this is true even if the categorical claim—that perception is *always* shaped by mental set—is proved false.)[4]

Most of the relevant experimental cases adduced by the New Look psychologists aim to show a difference in the perceptual experience had by experimental subjects

[3] A number of critics have questioned Pylyshyn's conclusions in this general way (Bermudez 1999; Macpherson, forthcoming; Moore 1999; Noë and Thompson 1999). It is also worth noting that Pylyshyn's empirical claim can be challenged (see Boynton 2005; Kamitani and Tong 2005).
[4] See Balcetis and Dunning 2006 and van Ulzen 2008 for brief historical discussions of the rise and fall of the New Look movement, as well as new studies in the New Look spirit.

(e.g. seeing a stimulus as bigger than it is), where this difference is supposed to be a direct causal result of cognitive states of the perceiver.[5] Critics of the New Look movement have most commonly deflected these putative cases of cognitive penetration by the following three strategies.

First, critics claim that what is affected by the subject's cognitive states is the subject's memory rather than her perceptual experience. Subjects recall the stimulus to be some way (i.e. a way different than if the relevant cognitive states are controlled for) as a result of some other cognitive state, and report a memory of the stimulus rather than a perceptual experience. This evidences cognitive penetration of cognition. And this is uncontroversial: memories can be faulty, and in ways influenced by what we believe, desire, or otherwise think. Call this the *memory interpretation*.[6]

A second strategy is the *attention-shift interpretation*. This interpretation maintains that in the cases in question, cognitive states of the experimental subjects cause a shift in attention, generally involving some overt action, which then results in the change in perceptual experience. Thus the link between cognition and perception is mediated by an external action. This is no different in kind, critics urge, from an ordinary perceptual scenario where one, for example, has some belief about one's environment (e.g. there is an irritating noise somewhere around here) and this belief causes some action (e.g. going to look for the thing making the irritating noise), which in turn results in a changed perceptual experience (e.g. seeing the squeaky faucet). This

[5] Bruner and Goodman 1947 is discussed in detail below. See also Postman, Bruner, and McGinnies 1948; Bruner, Postman, and Rodrigues 1951; Bruner and Minturn 1955.
[6] For one example, see McCurdy 1956.

familiar cognitive-behavioural dynamic is important to everyday life, but unless cognitive penetration is trivially rampant, this is not cognitive penetration. Cases involving shifts of attention, this interpretation suggests, are to be treated similarly: these scenarios lack an appropriate internal connection, and so there is nothing in this causal chain to properly call 'penetration' (see Pylyshyn 1999: 343).[7]

Finally, critics have suggested that the experimental subjects are not reporting a cognitively affected perceptual experience, but instead a judgement of the perceived stimulus. So the perceptual experience of the stimulus remains unaffected. At most, the subject judges or evaluates the stimulus in a way she would not if she lacked some background cognitive state/s. This difference manifests in the different reports of the experimental subjects versus the control subjects in the New Look studies. Call this the *judgement interpretation*.[8]

Grant that if any putative case of cognitive penetration can be interpreted in one of these alternative ways, then the critics are correct: it is *not* a genuine case of cognitive penetration of experience. We can then define cognitive penetration so as to rule out these interpretations, and ask if any case plausibly meets the definition. If the answer is 'yes', then the critics must secure some alternative interpretation to deflect the case/s. Here is such a definition:

[7] Fodor also appeals to this general response in his debate with Paul Churchland on the theory-ladenness of perception/observation (see Fodor 1988; Churchland 1988; see also Fodor 1983).

[8] For extended discussion of these and other strategies for the cognitive impenetrability theorist, see Macpherson, forthcoming; AUTHOR MASK B-XXXX.

(CP) A perceptual experience E is cognitively penetrated if and only if (1) E is causally dependent upon some cognitive state C and (2) the causal link between E and C is internal and mental.

The definition requires a few qualifications. First, assume an orthodox understanding of 'cognitive state.' Cognitive states are representational and possess some kind of linguistic or propositional content, and they play a role in reasoning and decision making. On this minimal understanding of 'cognitive state', standard cognitive states include beliefs, desires, values and, perhaps, emotions.

Second, clause (2) says that if an unscreened internal cause involves a cognitive state—that is, the causal chain runs from experience back to a belief, desire, or some other cognitive state without deviating from internal mental processes—then the perception depends (internally) upon a cognitive state. Counterfactually, had C not been present in the process, E would not be had by the subject. C is thus a necessary causal condition for E. Understood probabilistically, C is not a strictly necessary causal element, but one that is highly relevant to the probability of that perceptual experience; E is more likely to be had when C is present, and less when not present. The preferred notion of causation is of little matter so long as the internal causal dependence is maintained.[9]

Finally, (CP) excludes obviously non-genuine cases of cognitive penetration. For example, a desire to see the show, coupled with a true belief about the location of the

[9] One should also note that C is a non-sufficient cause of E. There are other relevant causal factors. Again, counterfactually, cognitive states are causally relevant, such that if state C had not been present in the perceptual-cognitive system of the agent, then that agent would not have had perceptual experience E.

show, may result in a perceptual experience of the show (several experiences in fact). But this should not count as an instance of cognitive penetration of experience, else the concept 'cognitive penetration' becomes trivial. (CP) delivers the appropriate result: such a case fails to satisfy our definition, since it fails to satisfy clause (2). In cases like this, a cognitive state (or some cognitive states) motivates an action (or set of actions) which eventually results in the relevant experience. The perceptual experience thus causally depends upon the relevant cognitive state/s. Clause (2) insures, however, that such cases are not instances of cognitive penetration, since in each case the cognitive states are screened from being *internally*, causally efficacious: the cognitive states cause an (external) action which eventually results in the experience.[10]

Now for the importance of (CP): A perceptual experience that satisfies this definition cannot be interpreted in any of the three ways described above. Clause (1) of (CP) rules out both the memory and judgement interpretation, since it requires a cognitive influence on perception, rather than just an influence on some other cognitive state in the system. Clause (2) of (CP) rules out the attention-shift interpretation, since it requires a non-externally mediated causal link between the cognitive state and the perceptual experience. The question now becomes: are there any experimental cases that satisfy (CP)? The answer is 'yes'. In fact, one of the most compelling and instructive cases comes from an early New Look study.

In a now famous experiment, Jerome Bruner and C.C. Goodman tested perceptual experiences of objects of social value (Bruner and Goodman 1947). Three groups (10

[10] In this way, CP is consistent with other recently proposed definitions of cognitive penetrability, discussed above (Macpherson, forthcoming; Siegel, forthcoming A).

persons per group) of 10 year old children, two experimental and one control, were put before a wooden box with a glass screen on its face. In the centre of the screen was a small patch of light, nearly circular in shape, the circumference of which could be adjusted by a small knob located on the bottom right corner of the box. The two experimental groups of children were presented with ordinary coins of varying values. As they looked at the coins, placed flat in the palm of the left hand, positioned at the same height and six inches to the left of the adjustable patch of light, they were asked to adjust the patch to match the size of the presented coin. The subjects could take as much time as they liked to complete the task. The control group was instead presented with cardboard discs of sizes identical to the relevant coins, and asked to perform the same task. In the experimental group, perceptual experience of the coins was "accentuated." The experimental subjects systematically overestimated the size of the coin, and sometimes by a difference as high as 30% as compared with control subjects.[11]

The second experimental variation divided experimental groups into subgroups comprising "rich" and "poor" children. The task was the same, except only real coins were used. Here, rich children, as the previous results would suggest, still overestimate the size of the coins, but at percentages significantly lower than the poor children. Indeed, poor children systematically overestimate the size of coins, by as much as 50%, and by differences as high as 30% as compared to rich children.[12]

---

[11] For example, experimental subjects overestimated the size of a dime by an average of 29%; controls underestimated the size of the cardboard analogue of a dime by -1%.

[12] Although there were a number of theorists critical of particular details and the broad scope of the New Look approach (Klein, Schlesigner, and Meister 1951; Carter and Schooler 1949; Lysak and Gilchrist 1955), the Bruner and Goodman 1947 results have been broadly replicated by a number of similar studies at least insofar as these studies all evidence some higher level effect on perceptual experience (in addition to those listed above

Bruner and Goodman's explanation was that the social value or desire for money was somehow affecting the perceptual experiences of the children. And indeed this case *prima facie* satisfies (CP). The experimental subjects have a perceptual experience, the character or content of which causally depends on a cognitive state, in this case, a desire or value. And the causal link between the experience and cognitive state is internal and mental.

Nonetheless, critics have traditionally rejected this as a genuine case of cognitive penetration, and by appeal to one of the strategies outlined above. However, this rejection is a mistake premised on a failure to carefully consider Bruner and Goodman's experimental procedures. In all the variations described above, subjects took as much time as they needed to adjust the light patch to match the size of the coins. The coins were presented at the same time as, at the same horizontal level as, and six inches to the left of, the adjustable light patch. Subjects did *not* visually inspect the coin and then shift to a distinct visual field, adjusting the light patch by memory. Instead, their task was to adjust the patch of light to match what they were seeing, *while* they were seeing it. In no relevant sense were they forced to base their report just on memory. The memory-interpretation thus fails.

For the same reason, the attention-shift interpretation fails: subjects did *not* attend to one stimulus (the disc or coin) and then shift attention to a distinct visual field where the second stimulus (the adjustable light patch) was located. The subjects would have shifted their gaze from disc/coin to light patch, but this slight shift would

---

in fn 5, see Bruner and Postman 1948; Dukes and Bevan 1952; Bruner and Rodrigues 1953; Blum 1957; Holzkamp and Perlwitz 1966). Only this last result is needed for present purposes; no commitment to the full scope and detail of the New Look theory is required.

not have differed across control and experimental subjects (and so would fail to explain the relevant differences between the two groups).

The judgement interpretation is most commonly used to deflect cases like the Bruner and Goodman case. The only version of this interpretation that is inconsistent with cognitive penetrability is one that claims that the perceptual experiences of the subjects are accurate across control and experimental subjects alike, while the experimental subjects make a misjudgement of the size of the coins. This interpretation is less plausible than the interpretation it opposes, since it requires attributing a judgement or belief to the child which does *not* correspond to the perceptual experience that she has simultaneously with that judgement or belief. It requires that the subject, while inspecting the visual stimulus—which is, again, at a constant location six inches to the immediate left of the adjustable light patch—consistently makes erroneous judgements about what she is seeing.

Here, the cognitive impenetrability theorist might respond by invoking instances where perception and judgment do come apart in just this way. So, for example, although one sees the Müller-Lyer lines as being of different lengths, one believes (if one knows the illusion) that the lines are of the same length. And indeed one cannot manage to see them accurately in spite of this background knowledge (Fodor 1983, 1985, 1988; Pylyshyn 1999). So, the critic would argue, a consistent mismatch between simultaneous experience and judgment is not so uncommon, and perhaps Bruner and Goodman's subjects can be explained similarly. However, the subjects in the Bruner and Goodman experiments are importantly different from standard perceivers of the Müller-Lyer and other such illusions. When one judges and reports that the Müller-Lyer lines are of the same length, one bases this report *not* on perceptual experience, but on knowledge of the illusion. Bruner and Goodman's subjects are different in this regard: they intend for their report to be one of what they see. If asked, the subjects would certainly report that they are matching the light patch to what they are seeing. To treat these subjects like perceivers of the Müller-Lyer illusion requires that they are systematically mistaken about this: the subjects are not correctly reporting what they see.

The judgement interpretation, then, must maintain that these subjects are continually ignoring, remaining unconscious of, or somehow otherwise failing to accurately report their perceptual experience. This, given the experimental circumstances, is far less plausible than the interpretation it opposes—namely, that experience itself is penetrated and then reported by the subjects. The judgement interpretation thus fails.[13]

If the above discussion is successful, then the standard strategies for dismissing putative cases of cognitive penetration fail to deflect Bruner and Goodman's results as genuine evidence for the cognitive penetration of perception. And we should note that there is flexibility regarding how these kinds of cases are further interpreted, in the following respect. We favour interpreting the results so the cognitively influencing state is a desire or value. Alternatively, one might think that it is a concept or belief regarding money that is causally operative. Or one might favour a less direct mechanism, where a belief or concept influences a visual image, which in turn

[13] There are other experimental cases that seem to evade the same strategies. Fiona Macpherson argues that Delk and Filenbaum's 1965 colour perception results, among others, cannot be deflected by the standard strategies (see Macpherson, forthcoming). For recent empirical studies, see Balcetis and Dunning 2006; Balcetis and Dunning, 2010; van Ulzen 2008.

influences the character of experience.[14] In any case, encapsulated modularity will be challenged by this case (and others), since whether it is a desire, value, belief, or some other higher-level mental state, there is *some* non-perceptual state (internally) influencing experience. And therefore, as we will now argue, perception is not operating in an informationally encapsulated way.

Recall that perceptual systems are paradigms for modular systems. And recall further that if one is an encapsulated modularity theorist, then one commits to the informational encapsulation of modules. The informational encapsulation of perceptual modules entails cognitive impenetrability. Finally, the cognitive penetration of perceptual experience entails, at some level, the cognitive penetration of perceptual processing. Thus any legitimate case of the cognitive penetration of experience undermines the purported informational encapsulation of the relevant perceptual systems, and in turn challenges any modularity theorist that commits to this feature as necessary for modules.

Here, finally, is the first horn of the dilemma for modular architectures of the mind. There are legitimate cases of the cognitive penetration of experience. Bruner and Goodman 1947 is the example we have defended; but many of the others mentioned above, old and new, could be defended in similar fashion.[15] And so perceptual systems—in this case *vision*—are not informationally encapsulated. Any modularity theory that commits to informational encapsulation (and by implication: cognitive impenetrability) as necessary for modularity is therefore threatened.

[14] See Macpherson (forthcoming) for such an account.

[15] See Siegel (forthcoming A; forthcoming B), Macpherson (forthcoming), and AUTHOR MASK B-XXXX, for further analysis of relevant empirical studies.

To clarify our critique, it will be useful to briefly consider a hypothetical defence for the encapsulated modularity theorist in response to this first horn of the dilemma. Our suggestion is not that perception (or vision, more specifically) is, as it were, unencapsulated through-and-through. As discussed above, the entailment relations between the cognitive penetration of perceptual experience and the cognitive penetration of perceptual processing would not support this last inference. So, the modularity theorist might retort, the penetration of experience is compatible with the impenetrability (and thus encapsulation) of *some* (but not all) components or systems in perceptual processing. Some components of perceptual systems may be strongly modular.[16]

All of this may be correct, but it does nothing to save a modular *architecture* of perception. For example, feature detecting components like groups of simple and complex cells in the primary visual cortex are likely encapsulated, as are many other neural circuits and low-level components in the overall visual system. In fact, theorists may be able to maintain that certain sub-systems in vision—for example, Pylyshyn's early vision—are encapsulated in spite of the penetration of visual experience. This would be to maintain the commitment to informational encapsulation and thus a *strong* notion of 'module'. But note that the *scope* of modularity on such a view is significantly weakened: such a modularity theorist can only claim that *some* of the visual system is modular and, importantly, cannot claim that vision is, generally, modular. This last claim *is* inconsistent with genuine cases of cognitive penetration.

[16] Thanks to XXXX for pressing us to consider this response.

It is important, moreover, not to overemphasize a characterization of (perceptual) modularity as concerning only computational processes. This is for the simple reason that an interest in mental architecture is guided not merely by goals of psychological modelling. Another crucial issue of relevance is epistemic: a modular theory of perception promises a preferable epistemology, one where perceptual systems rapidly deliver perceptual representations in a way not prone (or less prone) to errors introduced by the cognitive agent. As Fodor puts the point, the "function of perception is to deliver to thought a *representation* of the world." And since here the goal is to represent "[n]ot the distant past, not the distant future and not…what is very far away…it is understandable that *perception* should be performed by fast, mandatory, encapsulated, etc. systems…" (Fodor 1985: 5; emphasis added). The systems in question are sub-personal modules, but the representations they provide or give rise to are personal-level experiences. Given the epistemic role that such representations are supposed to serve, and the putative epistemic advantage of modular perceptual systems, the modular theorist should be no happier with evidence for penetrated experience than he is with evidence for unencapsulated perceptual processing.[17]

Where does this leave the view? The claim that some individual low-level circuits are encapsulated and thus strongly modular is largely uncontroversial among cognitive scientists. And the claim that some sub-systems—even Pylyshyn's early vision—in perception are strongly modular is insufficient to support the claim that

[17] Fodor makes similar suggestions elsewhere; see, for example, his discussion of "perceptual identifications" (1983: 68-71). And Pylyshyn (1980) makes similar commitments, claiming that the reliability of perception requires cognitive impenetrability. For further discussion of the epistemic consequences of cognitive penetrability, see Lyons, forthcoming.

the general structure of perception (or, more specifically, vision) is strongly modular. In short, one cannot save an encapsulated modularity by appeal to encapsulated components or sub-systems of perception. To do so would be to opt for strength of modules over scope, in turn undermining the theory as an *architecture* of the mind.

II. Informationally unencapsulated modules: A challenge for the massive modularity hypothesis

A number of recent theorists have weakened the notion of modularity with respect to Fodor's original characterization and, in particular, with respect to informational encapsulation. This change in the notion of modularity tends to accompany a broadening of the scope of modular theories. Thus, massive modularity theorists take much if not the whole of the human mind to be modular, including higher level conceptual and cognitive systems. If, as we have argued in the previous section, cognitive impenetrability (and thus encapsulation) seems too strict a requirement on the modularity of perceptual systems, then it makes sense not to require it of higher-level cognitive systems. Weakening modularity in this way, however, comes with significant costs to any modular account of cognition. First, it weakens the explanatory value of modular architectures. Second, it threatens the internal coherence of modular theories.

Peter Carruthers, a massive modularity theorist, argues that

> *if a thesis of massive mental modularity is to be even remotely plausible*, then by 'module' we cannot mean 'Fodor-module'. In particular, the properties of having proprietary transducers, shallow outputs, fast processing, significant innateness or innate channelling, and *encapsulation* will very likely have to be struck out. (Carruthers 2006: 12; emphasis added.)

According to Carruthers, massive modularists should expect most (if not all) central cognitive modules to be *un*encapsulated. He writes:

> …even where a system has been designed to focus on and process a particular domain of inputs, one might expect that in the course of its normal processing it might need to query a range of other systems for information of other sorts. (Carruthers 2006: 10).

In other words, an unencapsulated module, in order to perform its task, will often need to compute over information that is made available by other systems. For example, the mind-reading system "may need to query a whole range of other systems for information relevant to solving the task in hand" (Carruthers 2006: 11).

Evolutionary psychologists, many of whom subscribe to the massive modularity hypothesis, also tend to argue for (or assume) the compatibility of modularity with unencapsulation. Hagen (2005) explicitly states what is often implicitly assumed in this field:

> Why, except when processing speed or perhaps robustness is exceptionally important, should modules not have access to data in other modules? Most modules should communicate readily with numerous (though by no means all) other modules when performing their functions, including querying the databases of selected modules (163).

Any such modularity theorist thus claims that systems, like the mind-reading system, can be modular *in spite of* being informationally unencapsulated. As Carruthers suggests, this might be a necessary adjustment of a general modular architecture for the simple reason that anything stronger is implausible.

One main theoretical advantage of, and indeed motivation for, postulating modular architectures is that they explain dissociations between perceptual and cognitive functions. A cognitive task *A* is said to be dissociated from cognitive task *B* when at least some individuals are observed who show a significant deficit with respect to *A* in the absence of a corresponding deficit in *B*. *A* and *B* are said to be *doubly* dissociated when, in addition, we observe individuals in whom *B* is significantly impaired without a corresponding deficit in *A*. Cognitive scientists generally hold that dissociations are signs of functional independence, and will often postulate the existence of cognitive modules on the basis of these behavioural patterns. If *A* is observed to fail when *B* does not, then one may infer that *A* involves a mechanism, or module, *M* that *B* does not recruit. When *M* is obstructed, it is argued, *A* fails and *B* does not. In the case of a double dissociation, the inference is stronger, namely that *A* and *B* are each performed by a mechanism, or module, that the other does not recruit.

A classic, although by no means uncontroversial, case of the stronger version of the inference concerns the face recognition module hypothesis. Many theorists postulate the existence of a face recognition module on the basis of a double dissociation between face recognition and other cognitive and perceptual functions such as visual object recognition. On the one hand, there are well-documented cases in the neuropsychological literature of patients who have a face recognition deficit (prosopagnosia) but who appear to retain all other cognitive and perceptual capacities, including the ability to respond to and identify an array of physical objects (i.e. non-faces) (Rossion et al. 2003). On the other hand, there are reports of patients with impaired visual object recognition who appear to retain all other cognitive and perceptual capacities, including the ability to respond to and identify familiar faces (Rumiati & Humphreys 1997). This is, ostensibly, a double dissociation between face and object recognition and it has motivated some theorists to posit two corresponding, separate modules: a face recognition module and a visual object recognition module (Coltheart 1999: 119).

This reasoning from dissociation data to modularity—call it the *functional modularity inference*—has been central to the development of modern neuropsychology. In the last thirty years, philosophers and cognitive scientists have refined concepts of dissociation and narrowed the scope of the inference, and there is an emerging consensus that the inference should be understood as an inference to the best explanation, where one infers that a cognitive system is modular on the grounds that this hypothesis best explains a set of dissociation data, given the fact that modular systems would produce dissociations if damaged in different ways (Shallice

1988, Coltheart 2001).[18] The status of the inference as a central methodological tool, however, is very much a matter of debate. Various authors have argued, on both theoretical and empirical grounds, that the existence of a double dissociation between subjects' performance on two different cognitive tasks does not necessarily constitute strong evidence for the existence of separate cognitive functions or modules (Dunn & Kirsner 2003; Juola, & Plunkett 2000; Plaut 1995; Van Orden, Pennington, & Stone 2001). We take no side in this debate. Instead, we question whether an inference to the best explanation is supported *when modules are assumed to be unencapsulated*.

Let us suppose, then, as the weakened modularity theory we're considering does, that modules are *not* encapsulated. Suppose, for example, that the alleged face and object recognition modules are not encapsulated, that they both often need to compute over information made available by other systems in order to perform their tasks. This means that a double dissociation between face and object recognition could occur *even if* both alleged modules remained intact (i.e. were not damaged). This would occur, for instance, if both modules need to compute over information normally made available by other systems and damage to these other systems (or damage to some pathway between them and the two modules) prevents the availability of the needed information. In this case, the double dissociation could no longer be taken as a sign that the two modules are functionally independent, as both of them could fail to perform their tasks for reasons that have nothing to do with a failure of their respective mechanisms. But since it is the assumed functional independence of cognitive modules, in the functional modularity inference, that is supposed to explain

[18] See Shallice 1988 for a detailed discussion of this methodology.

the existence of dissociation data, it is therefore difficult to see how the face and object recognition module hypotheses could, in this case, *best* explain the observed double dissociation between face and object recognition.[19]

By contrast, a double dissociation between face and object recognition is adequately explained by encapsulated modularity. Importantly, because the face and object recognition modules would in this case be informationally encapsulated, their normal functioning would thereby not depend on information made available by other systems. A double dissociation between face and object recognition would thus suggest that the modules themselves have been separately damaged.[20]

This point about the explanatory weakness of unencapsulated modularity is worth further emphasis, since it suggests the important distinction between dissociation at the level of behaviour (behavioural dissociations) and dissociation at the level of mechanisms (mechanistic dissociations). As explained above, unencapsulated modules can give rise to *behavioural* dissociations when there is damage to the systems that provide these modules with information required for task performance. But the fact that unencapsulated modules can give rise to behavioural dissociations does not imply that the mechanisms underlying their function are dissociable.[21] In fact, there is

[19] The putative theory of mind module is another case, one where unencapsulated modularity would fail to explain behavioural dissociations between mind reading and other cognitive capacities. See Gerrans and Stone (2008) for an extended discussion of this case.

[20] This is not to say, however, that such an encapsulated modular architecture is the only plausible explanation of the double dissociation between face and object recognition. Even with encapsulated modularity, the functional modularity inference remains abductive.

[21] The distinction between behavioural dissociation and mechanistic dissociation mirrors, though imperfectly, the distinction between perceptual experience and perceptual processing. And just as one must be careful in making inferences from (putative) facts about experience to (putative) facts about processing, and vice versa, one must be careful about making inferences from facts about behaviour to facts about mechanism, and vice versa.

conceptual tension between the notions of unencapsulation and dissociation at the level of mechanisms.

To see this, consider a minimal conception of modularity. As Carruthers suggests, a module must be, at the very least, a "dissociable functional component" (Carruthers 2006: 2). This minimal conception of modularity is what gives the functional modularity inference its theoretical force. Behavioural dissociations are explained by modular architectures—and are thus signs of functional independence between cognitive mechanisms—since dissociability (at the level of mechanisms) *implies* functional independence and modules are dissociable systems. We agree that a minimal notion of modularity should include dissociability, since without it a modular architecture would reduce to functional decomposition. And functional decomposition—understanding the mind in terms of functional components and sub-components—is uncontroversial as an approach, except perhaps in some connectionist quarters.

On the one hand, therefore, a cognitive system *S* is considered functionally independent from another system *O* if *S* and *O*'s function can be dissociated. This means that *S*'s function is not affected by what happens to *O*, and that *S* can be modified (or damaged) without affecting *O*'s function.

On the other hand, *S* is considered unencapsulated relative to *O* if *S* needs to compute over information made available by *O* in order to perform its task. This means that *S* depends on *O* for its normal functioning. *S* is thus functionally dependent on *O*, and is therefore *not* dissociable from *O* (since dissociability implies functional independence). In sum, if functional independence is understood in terms of dissociation, as the minimal conception of modularity suggests, then *S* cannot both depend on information provided by *O* and be dissociable from *O*.[22]

An illustration may help. Carruthers explains the minimal conception of modularity with the following analogy.

> The hi-fi is modular if one can purchase the speakers independently of the tape-deck, say, or substitute one set of speakers for another for use with the same tape-deck. Moreover, it counts towards the modularity of the system if one doesn't have to buy a tape-deck at all—just purchasing a CD player along with the rest—or if the tape-deck can be broken *while the remainder of the system continues to operate normally*. (Carruthers 2006: 2; emphasis added).

Carruthers goes on to suggest that although operationally distinct in the above ways, the components of the hi-fi, once conjoined as a system, do depend upon one another in other ways: the CD player requires the amplifier to distribute sound, the speakers require input from the amplifier to make sounds, etc. Indeed, some of these dependence relations will be asymmetric: the CD player needs the amplifier to distribute sound, but not vice versa. The important point to note for present purposes is that in spite of these dependence relations, the hi-fi components are (relevantly) computationally autonomous: the CD player may require the amp to deliver its output, but it does not need to compute over information made available

[22] Note that the claim is not that a cognitive system *S* cannot be both unencapsulated relative to another system *O* and dissociable from *O simpliciter*. *S* can in fact be both unencapsulated relative to *O* and dissociable from *O*, as long as *S* does not need to compute over information provided by *O* in order to perform its task. Our claim, therefore, applies to cases in which *S* does need to compute over information made available by *O*, cases which according to massive modularity theorists should be very common.

by the amp in reading data off of a CD. In other words, in performing its task, it is encapsulated from the amp, the speakers, and so on. Likewise for other components in the system: the tape-deck reads data, the tuner acquires a radio signal, the speakers deliver a range of sounds, and so on, all independently.

Thus the hi-fi analogy is a useful one, at least for modularity as traditionally understood. The trouble is that unencapsulated cognitive modules are relevantly disanalogous to hi-fi components. Like the hi-fi modules, the cognitive modules envisaged by massive modularity work together, exchanging input and output, and often asymmetrically. But unlike the hi-fi components, an unencapsulated module *M*, as per the massive modularity theorist, will often depend for its normal operation on other components in the system. And this means that *M* will not, in these cases, be dissociable from these other components. Here again *M* cannot both depend on information provided by other systems and be dissociable from them.

All of this challenges recent theories that expand the scope of modularity by weakening the notion of modularity so as not to require informational encapsulation; this is the second horn of our proposed dilemma. First, weakening modularity to this degree weakens the explanatory value of modular architectures, which in turn weakens the functional modularity inference. Second, the very notion of unencapsulated module appears to be at odds with the core of modularity: viewing modules as dissociable functional components.

III. Functional independence without modularity

Over the past century and a half, a large body of neuropsychological data, primarily in the form of dissociation data, indicates that there are specialized neural circuits in the brain and that there is structure-function correspondence between these circuits and patterns of behaviour.[23] In fact, modular and non-modular theorists alike see specialization within the brain as an undisputed fact. Both sides, therefore, would agree in substance with Norman Geschwind's account of the general architecture of the brain as "more or less specialized groups of cells connected by relatively discrete pathways" (Geschwind 1965). What is at issue is how best to characterize this specialization.

Recall that a module can be minimally characterized as a dissociable functional component. We have argued that there is tension between this conception of modularity and the idea that modules often need to compute over information made available by other systems in order to perform their tasks. The argument depends on understanding functional independence in terms of dissociability: a cognitive system *S* is functionally independent from another system *O* if *S*'s function can be dissociated from *O*'s function. Clearly, dissociability implies functional independence, but it is now worth asking whether functional independence implies dissociability. We now argue that functional independence can be defined in a way that does not imply dissociability, thereby allowing for cognitive systems that are functionally independent *and* both unencapsulated and non-dissociable. We argue further that the resulting notion of a functionally independent system, while considerably weaker

[23] More recently, functional neuroimaging data (e.g. fMRI, PET), in the form of selective activations of brain areas for certain tasks, also point to the wide range of specialized neural circuits and structure-function correspondences in the brain, although the methodology in this case differs from the standard behavioural dissociation logic in neuropsychology.

than the minimal conception of a module, is sufficiently strong to characterize the structure-function correspondence between brain structures and particular perceptual and cognitive functions.

To see how a system *S* can be functionally independent from another system *O* without being dissociable nor encapsulated from *O*, consider the case of Broca's area (BA). This area is involved in language processing (production and perception of speech), and in order to contribute to this cognitive capacity it needs to compute over information made available by other areas, one of which is the superior temporal sulcus (STS) that contains phonological representations (Hickok & Poeppel 2007). BA is thus unencapsulated relative to the STS. Moreover, since BA's contribution to language processing would be affected if the STS were damaged, it is therefore not dissociable from the STS. Nevertheless, there is a sense in which BA is functionally independent from the STS.

To see this, let us first introduce a distinction between the low-level computational operations, or "workings", performed by BA, and the higher-level cognitive "uses" to which it is put (AUTHOR MASK A-XXXX, Anderson, 2010). We know that BA is put to a number of linguistic and non-linguistic uses—for example, it is involved in both musical and linguistic syntactic processing, in object manipulation, and in action sequencing and action perception (Nishitani et al. 2005; Patel 2003; Schubotz & Fiebach 2006). This, in turn, has been interpreted as evidence that BA's contribution to these various cognitive uses could be performed by a "reusable" set of low-level computational operations, or workings—e.g. hypersequencing operations in the form of "detection, extraction, and/or representation of regular, rule-based patterns in temporally extended events" (Schubotz & Fiebach 2006: 501).[24] 'BA's function', therefore, can be interpreted in two different ways depending on whether one is referring to its local workings or to its higher-level cognitive uses.

In the light of this distinction, we can now specify the sense in which BA is functionally independent from the STS. BA is functionally independent from the STS *with respect to hypersequencing operations* in the sense that BA performs these operations and has the capacity to perform them even if the STS failed to compute anything. Generalizing:

> (FI) A system *S* is *functionally independent* from another system *O with respect to working W*, iff *S* performs *W* and has the capacity to perform *W* even if *O* failed to compute anything.[25]

And accordingly:

> (FI system) A system *S* is an *FI system with respect to working W*, iff *S* performs *W* and has the capacity to perform *W* even if no other systems computed anything.

A few points are in order. First, to say that *S* has the *capacity* to *perform W* even if another system *O* failed to compute anything is to say that *S* possesses the right kind

[24] See Fiebach and Schubotz (2006) and Fedorenko et al. (2009) for other proposals of BA's workings.

[25] This formulation is adapted from the analysis of isolability provided by Lyons (2001): "A substrate *S* is *isolable with respect to task T* iff *S* performs task *T* and could do so even if nothing else computed any (cognitive) functions"(289).

of machinery to perform *W* given that it is provided with the right kind of information. The second conjunct of the condition for functional independence must thus be read counterfactually.

In the case in which *S* performs *W* over information provided by *O*, this means that *S* has the right kind of machinery to perform *W* over the kind of information that *O* normally provides. For example, if we assume that BA does in fact perform sequencing operations over phonological representations made available by the STS, then to say that BA has the capacity to perform sequencing operations even if the STS failed to compute anything (and thus failed to make available the relevant information) is to say that BA has the right kind of machinery to perform these operations over the kind of information that the STS normally provides. (In this case, it is easy to see that BA does have this capacity since, as noted, BA's workings appear to be put to a wide range of cognitive uses in which the STS does not participate.)

Second, the notion of FI system is considerably weaker than the *minimal* notion of a module as a dissociable functional component. FI does not require systems to be functionally dissociable (as explained above). In fact, the basic cognitive architecture suggested by FI—the idea that the mind is composed of FI systems—cannot, in any non-trivial way, count as modular. To see this, suppose that modules are no more than FI systems or networks of FI systems. Suppose, for example, that BA plus the STS, plus some other systems constitute the speech production module. This is to *assemble* a module out of FI systems: decomposing the mind by tasks or capacities and attributing these capacities to an assemblage of systems. By this method, any identifiable cognitive capacity could be turned, trivially, into a module, since any identifiable cognitive capacity could, under this architecture, be implemented by an FI system or a network of FI systems. What is trivial is not that our proposed cognitive architecture could potentially apply to any identifiable cognitive capacity: this is exactly what a basic functional architecture should do. What's trivial, instead, is calling any FI system or network of FI systems capable of implementing a cognitive capacity a 'module'; labeling each cognitive capacity a 'module' adds nothing of theoretical import to the functional decomposition approach to studying the mind.

There is, in addition, another aspect of FI that makes it weaker than the minimal conception of modularity. In order for a system *S* to be functionally independent from another system *O* (according to FI), the requisite condition applies only to one of the two levels of functional specification—i.e. it applies to workings (at the local level), but *not* to cognitive uses (at the systemic level).

For example, assuming again that BA does in fact perform sequencing operations over phonological representations made available by the STS, BA is functionally independent from the STS *with respect to its sequencing operations* (local workings), but it is *not* functionally independent from the STS *with respect to speech production* since both areas are jointly put to this higher-level use. To see this, consider a modification of (FI) where the working/use distinction is not made. Replacing 'working' with the more general term 'function' we have:

> (FI*) A system *S* is *functionally independent* from another system *O with respect to function F*, iff *S* has *F* and is capable of *F*ing even if *O* failed to compute anything.

Substituting BA for *S* and STS for *O* we have:

> BA is *functionally independent* from the STS *with respect to function F*, iff BA has *F* and is capable of *F*ing even if the STS failed to compute anything.

In this case, the condition is satisfied if *F* refers to sequencing operations, but it is not satisfied if *F* refers to speech production.

Importantly, the distinction between these two levels of functional specification is rarely operative in modular theorizing. For example, neuropsychologists typically attribute cognitive *uses* (as opposed to workings) to brain areas on the basis of dissociation data. Thus, Broca's area has been characterized as a speech production module; the temporoparietal junction has been characterized as a theory of mind module; the fusiform face area (FFA) has been characterized as a face recognition module, and so on.

This, in fact, should not be surprising since the functional modularity inference rests on dissociation data that are derived from an analysis of the performance of brain-damaged patients on various cognitive tasks. As such, these data consist of the specification of the *behavioural consequences* of the (mal)workings of various cognitive components, which means that they will naturally be expressed in ways that capture one or more of the cognitive *uses* of these components within the larger cognitive economy.

Third, and most importantly, FI is sufficiently strong to characterize the structure-function correspondence between brain structures and particular perceptual and cognitive functions. Recall that it is these phenomena, evidenced by the large body of dissociation data, that have traditionally motivated (and currently motivate) the modular approach. While the application of the traditional dissociation logic in neuropsychology has typically involved characterizing the specialization of a particular brain area in terms of the area's cognitive *uses* (as explained above), FI architectures involve characterizing the specialization of a brain area in terms of the area's cognitive *workings*. Accordingly, FI architectures conceive of structure-function correspondence as structure-*working* correspondence as opposed to structure-*use* correspondence.

We therefore agree with massive modularity theorists with respect to what needs to be explained: structure-function correspondences evidenced by dissociation data. We think, however, that unencapsulated modularity is inadequate as a basis for modular architectures of the mind, since it is both too weak to adequately explain what it is supposed to explain best, and is at odds with the core conception of modularity. Structure-function correspondences, we suggest, might be better explained by a different cognitive architecture, one that revolves around a non-modular notion of functionally independent system.

Acknowledgements

Bibliography

Anderson, M. (2010) 'Neural reuse: A fundamental organizational principle of the brain,' *Behavioural and Brain Sciences*: 33: 245-266.

Balcetis, E. and Dunning, D. (2006) 'See what you want to see: Motivational influences on visual perception,' *Journal of Personality and Social Psychology* 91(4): 612-25.

______(2010) 'Wishful Seeing: Desired objects are seen as closer' *Psychological Science*: 21: 147-52.

Bermudez, J. (1999) 'Cognitive impenetrability, phenomenology, and nonconceptual content,' *Behavioural and Brain Sciences* 22(3): 367-8.

Blum, A. (1957) 'The value factor in children's size perception, *Child Development* 28: 14–18.

Boynton, G. M. (2005) 'Imagining orientation selectivity: Decoding conscious perception in V1,' *Nature Neuroscience*, 8: 541–542.

Brewer, W. F., & Lambert, B. L. (2001) 'The Theory-Ladenness of Observation and the Theory-Ladenness of the Rest of the Scientific Process,' *Philosophy of Science* 68: 176-86.

Bruner, J.S., & Goodman C.C. (1947) 'Value and need as organizing factors in perception,' *Journal of Abnormal and Social Psychology* 42: 33-44.

Bruner, J. S., & Minturn, A. L. (1955) 'Perceptual identification and perceptual organization,' *Journal of General Psychology*, 21-28.

Bruner J. S., & Postman, L. (1948) 'Symbolic value as an organizing factor in perception,' *Journal of Social Psychology* 27: 203-208.

Bruner, J.S., Postman, L., and Rodrigues, J. (1951) 'Expectation and the Perception of Color' *American Journal of Psychology* 64: 216-227.

Bruner, J.S., & Rodrigues, J. S. (1953) ‘Some determinants of apparent size,’ *Journal of Abnormal and Social Psychology* 48: 17-24.

Carruthers, P. (2006) *The Architecture of the Mind*, Oxford: Oxford University Press.

Carter, L. F., & Schooler, K. (1949) ‘Value need and other factors in perception,’ *Psychological Review* 56,: 200-207.

Churchland, P. M. (1988) ‘Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor,’ *Philosophy of Science* 55: 167-187.

Coltheart, M. (1999) ‘Modularity and Cognition,’ *Trends in Cognitive Science* 3: 115-20.

Cowie, F. (2008) ‘Us, Them and It: Modules, Genes, Environments and Evolution,’ *Mind and Language* 23: 284-92.

Delk, J. L. and Fillenbaum, S. (1965) ‘Differences in Perceived Colour as a Function of Characteristic Color,’ *The American Journal of Psychology,* 78: 290-93.

Dukes, W. F., & Bevan, W. (1952) ‘Size estimation and monetary value: a correlation,’ *Journal of Psychology* 34: 43-53.

Dunn, J. C. & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex*, *39*(1), 1-7.

Farah, M. (1994) ‘Neuropsychological inference with an interactive brain: A critique of the locality assumption,’ *Behavioural and Brain Sciences* 17: 43-104.

Fedorenko, E., Patel, A., Casasanto, D., Winawer, J. and Gibson, T. (2009) Structural intergration in language and music: Evidence for a shared system. Memory and Cognition 37(1): 1-9.

Fiebach, C. J. & Schubotz, R. I. (2006). Dynamic anticipatory processing of hierarchical sequential events: a common role for Broca's area and ventral premotor cortex across domains? *Cortex*, *42*(4), 499-502.

Fodor, J. (1983) *Modularity of Mind*, Cambridge, MA: MIT Press.

______(1985) 'Precis of *The Modularity of Mind*', *The Behavioural and Brain Sciences* 8: 1-

5

______(1988) 'A reply to Churchland's "Perceptual plasticity and theoretical neutrality,"' *Philosophy of Science* 55: 188-19

Gentaz, E. and Rossetti, Y. (1999) 'Is haptic perception continuous with cognition?' *Behavioural and Brain Sciences* 22(3): 378-9.

Gerrans, P. and Stone, V. E. (2008) 'Generous or parsimonious cognitive architecture? Cognitive neuroscience and theory of mind'. *British Journal for the Philosophy of Science* 59: 121-41.

Geschwind, N. (1965). Disconnexion syndromes in animals and man. *Brain*, *88*, 237-294, 585-644.

Goodale, M. A. and Milner, D. (1992) Separate visual pathways for perception and action. *Trends in Neurosciences* 15(1): 20-25.

Hagen, Edward H. (2005). Controversial issues in evolutionary psychology. In D. M. Buss (Ed.), *The handbook of evolutionary psychology*. (pp. 5-67). Hoboken: John Wiley & Sons.

Hickok, G. and Poeppel, D. (2007). The cortical organization of speech processing. Nature Reviews Neuroscience, 8, 393-402.

Holzkamp, K. and Perlwitz, E. (1966) 'Absolute oder relative Größenakzentuierung? Eine experimentelle Studie zur sozialen Wahrnehmung,' *Zeitschrift für experimentelle und angewandte Psychologie* 13: 390-405.

Juola, P. & Plunkett, K. (2000). Why double dissociations don't mean much. In G. Cohen, R. A. Johnston, & K. Plunkett (Eds.), *Exploring cognition: damaged brains and neural networks*. (pp. 319-327). Psychology Press.

Kamitani, Y., & Tong, F. (2005) 'Decoding the visual and subjective contents of the human brain,' *Nature Neuroscience* 8: 679–685.

Klein, G. S., Schlesinger, H. J., & Meister, D. E. (1951) 'The effect of personal values on perception—an experimental critique,' *Psychological Review* 58, 96-112.

Lyons, J. (2001) 'Carving the Mind at its (Not Necessarily Modular) Joints,' *British Journal for the Philosophy of Science* 52 (2):277-302.

______(forthcoming) 'Circularity, Reliability, and Cognitive Penetrability of Perception',

Philosophical Perspectives.

Lysak, W., & Gilchrist, J. C. (1955) 'Value, equivocality, and goal availability as determinants of size judgments,' *Journal of Personality* 23: 500-501.

Machery, E. (2008), 'Massive Modularity and the Flexibility of Human Cognition,' *Mind and Language* 23: 263-72.

Macpherson, F. (forthcoming) 'Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism,' *Philosophy and*

Phenomenological Research.

McCurdy, H. G. (1956) 'Coin perception studies and the concept of schemata,' *Psychological Review* 63: 160–168.

Moore, C. (1999) 'Cognitive impenetrability of early vision does not imply cognitive impenetrability of perception,' *Behavioural and Brain Sciences* 22(3): 385-6.

Nishitani, N., Schürmann, M., Amunts K. and Hari, R. (2005) Broca's region: From action to language. *Physiology* 20:60–69.

Noë, A. and Thompson, E. (1999) 'Seeing beyond the modules toward the subject of perception,' *Behavioural and Brain Sciences* 22(3): 386-7.

Patel, A. D. (2003). Language, music, syntax and the brain. *Nat Neurosci*, *6*(7), 674-681.

Plaut, D. C. (1995). Double dissociation without modularity: evidence from connectionist neuropsychology. *Journal of Clinical & Experimental Neuropsychology*, *17*(2), 291-321.

Postman, L., Bruner, J. S., & McGinnies, E. (1948) 'Personal values as selective factors in perception,' *Journal of Abnormal and Social Psychology* 43: 142–154.

Pylyshyn, Z. (1980) 'Computation and cognition: issues in the foundations of cognitive science,' *Behavioural and Brain Sciences* 3:111-132.

______(1999) 'Is vision continuous with cognition? The case for cognitive impenetrability of visual perception,' *Behavioural and Brain Sciences* 22 (3):341-365.

Rossion, B., Caldara, R., Seghier, M., Schuller, A. M., Lazeyras, F., & Mayer, E. (2003).

A network of occipito-temporal face-facesensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain*, *126* (Pt 11), 2381-2395.

Rumiati, R. I. & Humphreys, G. W. (1997). Visual object agnosia without alexia or prosopagnosia. *Visual Cognition*, *4*, 207-217.

Shallice, T. (1988) *From Neuropsychology to Mental Structure*, Cambridge: Cambridge University Press.

Siegel, S. (forthcoming A) 'Cognitive Penetrability and Perceptual Justification,' *Nous*.

______(forthcoming B) 'The Epistemic Impact of Reasoning in the Basement,' *Philosophical Studies*

Sperber, D. (1994) 'The Modularity of Thought and the Epidemiology of Representations,' in L.A. Hirschfeld and S.A. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture*, New York: Cambridge University Press.

______(1996) *Explaining Culture: A Naturalistic Approach*, Oxford: Blackwell.

______(2001) 'Defending massive modularity,' In E. Dupoux (ed.) *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*, Cambridge, MA: MIT Press

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociation prove? *Cognitive Science*, *25*, 111-172.

van Ulzen, N.R., Semin, G.R., Oudejans, R., Beek, P. (2008) 'Affective stimulus properties influence size perception and the Ebbinghaus illusion', *Psychological Research* 72: 304–310.

Wilson, R. (2008), 'The Drink You Have When You're Not Having a Drink,' *Mind and Language* 23: 273-83.

Young, A. W. (1996). Face recognition. In J. G. Beaumont, P. M. Kenealy, & M. J. C. Rogers (Eds.), *The Blackwell dictionary of neuropsychology*. Blackwell.

AUTHOR MASK A-XXXX

AUTHOR MASK B-XXXX