

# What Are Modules and What Is Their Role in Development?

S. BUTTERFILL

---

**Abstract:** Modules are widely held to play a central role in explaining mental development and in accounts of the mind generally. But there is much disagreement about what modules are, which shows that we do not adequately understand modularity. This paper outlines a Fodoresque approach to understanding one type of modularity. It suggests that we can distinguish modular from nonmodular cognition by reference to the kinds of process involved, and that modular cognition differs from nonmodular forms of cognition in being a special kind of computational process. The paper concludes by considering implications for the role of modules in explaining mental development.

## 1. What Are Modules?

Jerry Fodor makes three claims about modules:

1. they are 'the psychological systems whose operations present the world to thought';
2. they 'constitute a natural kind'; and
3. there is 'a cluster of properties that they have in common ... [they are] domain-specific computational systems characterized by informational encapsulation, high-speed, restricted access, neural specificity, and the rest' (Fodor, 1983, p. 101).

Not all researchers agree about the properties of modules. That they are informationally encapsulated is denied by Dan Sperber and Deirdre Wilson (2002, p. 9), Simon Baron-Cohen (1995) and some evolutionary psychologists (Buller and Hardcastle, 2000, p. 309), whereas Scholl and Leslie claim that information encapsulation is the essence of modularity and that any other properties modules have follow from this one (1999b, p. 133; this also seems to fit what David Marr had in mind, e.g. Marr, 1982, pp. 100-1). According to Max Coltheart, the key to modularity is not information encapsulation but domain specificity; he suggests Fodor should have defined a module simply as 'a cognitive system whose application is domain specific'

---

Heartfelt thanks to Ian Apperly, Bill Brewer, John Campbell, David Charles, Naomi Eilan, Ross Harrison, Hemdat Lerman, Johannes Roessler, Jim Russell, Nick Shea, Matt Soteriou, Bart Streumer, and audiences in Laval, Reading, Sussex, Torino and York; most especially to Christoph Hoerl and several very generous anonymous referees whose advice and criticisms drastically improved this paper.

**Address for correspondence:** Department of Philosophy, University of Warwick, Coventry, CV4 7AL, UK.

**Email:** s.butterfill@warwick.ac.uk

(1999, p. 118). Peter Carruthers, on the other hand, denies that domain specificity is a feature of all modules (2006, p. 6). Fodor stipulated that modules are 'innately specified' (1983, pp. 37, 119), and some theorists assume that modules, if they exist, must be innate in the sense of being implemented by neural regions whose structures are genetically specified (e.g. de Haan, Humphreys and Johnson, 2002, p. 207; Tanaka and Gauthier, 1997, p. 85); others hold that innateness is 'orthogonal' to modularity (Karmiloff-Smith, 2006, p. 568). There is also debate over how to understand individual properties modules might have (e.g. Hirschfeld and Gelman, 1994 on the meanings of domain specificity; Samuels, 2004 on innateness).

In short, then, theorists invoke many different notions of modularity, each barely different from others. You might think this is just a terminological issue. I want to argue that there is a substantial problem: we currently lack any theoretically viable account of what modules are. The problem is not that 'module' is used to mean different things—after all, there might be different kinds of module. The problem is that none of its various meanings have been characterised rigorously enough. All of the theorists mentioned above except Fodor characterise notions of modularity by stipulating one or more properties their kind of module is supposed to have. This way of explicating notions of modularity fails to support principled ways of resolving controversy.

To illustrate, one person's claims about modularity might (a) exclude, (b) complement, or (c) re-formulate another's, yet it is often impossible to determine which of these is the case except by a stipulation. Take Elizabeth Spelke's notion of 'core knowledge systems' which resemble Fodor's modules in being domain-specific and informationally encapsulated (Spelke, 2000, p. 1233; 2003, p. 291) but may differ from them in other respects. The idea that there is such a thing as core knowledge has been proposed: (a) as excluding the claim that there are Fodorian modules, so that whereas you might have thought the mind contains Fodorian modules it actually contains core knowledge (e.g. Spelke, 1988); (b) as complementing this claim, so that core knowledge exists alongside Fodorian modules (e.g. Carey, 1995); and (c) as a re-formulation of this claim, perhaps broadening Fodor's notion of modularity (e.g. Hermer and Spelke, 1996). Which of these proposals correctly captures the relation between Spelke's core knowledge and Fodor's modules? As things stand there seems to be no way of answering this question except by stipulation.

The same problem occurs on a larger scale. Evolutionary psychologists typically think of modules as genetically specified and even as products of natural selection (Samuels, 1998, pp. 578–9) while sometimes denying information encapsulation (Buller and Hardcastle, 2000, p. 309). Developmental psychologists tend to emphasise domain specificity. And vision scientists are most likely to think modules are informationally encapsulated. How can we tell when these disciplines are tracking the several roles of a single module in evolution, development and cognition? How can we distinguish differences of emphasis or differences of subject matter from genuinely contradictory views involving modularity?

Defining notions of modularity by giving lists of characteristic properties does not enable us to resolve these kinds of question in a principled way. We need a

more rigorous approach to characterising modules, an approach that enables us to answer questions like ‘Are modules informationally encapsulated?’ without resorting to stipulation.

So what would be an adequate account of modularity? Interestingly, Fodor doesn’t define modules by specifying a cluster of properties (*pace* Sperber, 2001, p. 51); he mentions the properties only as a way of gesturing towards the phenomenon (Fodor, 1983, p. 37) and he also says that modules constitute a natural kind (see Fodor 1983: 101 quoted above). Suppose we have evidence that Fodor is correct to this extent: there are several psychological systems with some of the properties he supposes modules to have (they are domain-specific, informationally encapsulated and so on). Then we should ask, Why do some psychological systems have just this cluster of properties—why do these particular properties keep occurring together? Are these properties essential to the systems in question or could there be systems much like these that lack one or another of them? These are the sorts of question an adequate account of modularity would enable us to answer (as Fodor, 1983, pp. 37–8 says). Of course there may be more than one kind of module. But the same point applies: for each kind of module there is, we need to know not only which properties are common to all modules of this kind, but also why these systems have this particular cluster of properties. To borrow Fodor’s analogy with natural kinds, we need to identify the real essence of modularity.

By arguing that we lack a theoretically adequate account of modularity, I don’t mean to suggest that it isn’t already a useful concept. In my view, theorising about modularity is worthwhile only because there are cases where we are reasonably sure that hypotheses involving modules are explanatory.

In the next section of this paper I will offer a positive account of one kind of modularity. First, though, I want to continue arguing that existing accounts of modularity are all inadequate. Any reader already convinced that a better account of modularity is needed is invited to skip straight to Section Two.

There is a reason more basic than any given so far why current descriptions of modularity are inadequate. Each proceeds by identifying a cluster of properties, such as domain specificity and information encapsulation, and uses them to define modularity. What could any so-defined notion tell us about the architecture of real minds? In postulating modules we are not doing conceptual analysis or interpreting great scientists (or not only); we are aiming to understand something about how minds work. The claim that some cognition is modular is supposed to be a bold theoretical conjecture about how the mind is capable of generating predictions and insights. So modularity is a key explanatory notion. But no key explanatory notion can be adequately characterised by listing properties because the explanatory power of any notion depends in part on there being something which unifies its properties, and merely listing properties says nothing about why they cluster together. In particular, then, modularity cannot be adequately characterised by listing properties. This is the primary reason for wanting a better account of modularity.

You might object that modularity is not an explanatory notion, or that modularity is only explanatory because certain features of modules (such as information

encapsulation) are individually explanatory. To see that neither objection is correct we need to look at what modules explain.

Take Ken Cheng on spatial representation in rats. In a variety of tasks rats can make use of two types of cue, featural cues and geometric cues. Cheng (1986) showed that when a rat is lost (that is, uncertain of its location), it can use only geometric cues to determine its location. When the geometric cues available to the rat are consistent with it being in two or more places, the rat cannot use featural cues to find its bearings. So although rats normally use featural cues for a variety of purposes, lost rats cannot use featural cues to determine locations. To explain this surprising limitation, Cheng postulates that rats' navigational abilities depend on a module for spatial location which uses only geometric information (1986, p. 172). Cheng's findings have been challenged,<sup>1</sup> but I take it that Cheng was right at least insofar as the existence of such a module dependent exclusively on geometric information would explain his findings. However, this explanation would be inadequate if it were correct to characterise modules just by listing their properties. For what needs explaining is the rats' inability to use featural cues in certain situations. To say that this limitation exists because the relevant aspects of rats' cognition are informationally encapsulated is not to give an explanation which might easily be wrong but provides insight if correct; it's to do little more than label the findings. And this is what Cheng's explanation would amount to if there were nothing more to modularity than a cluster of properties. Since his explanation does amount to more than this, there must be more to modularity.

It is tempting to appeal to spatial metaphors in thinking through explanations like Cheng's. Just as academics tend to work at high-speed on domain-specific problems when they can cut themselves off from administrative centres, so we might attempt to explain the special properties of modules by saying that they are *cut off* from the central system. But it isn't clear how to turn this metaphor into an explanation. The spatial metaphor only gives us the illusion that we understand modularity.

Cheng's study of lost rats shows how modules are supposed to play a role in explaining cognitive organisation. They would not be able to play this role if we were to stipulate that 'module' means 'informationally encapsulated process.' This is why we should reject Greg Currie and Kim Sterelny's claim that 'encapsulation, being the essence of modularity, is what really makes a system modular and explains the other features.'<sup>2</sup> Even if information encapsulation could explain the other

---

<sup>1</sup> Nadel and Hupbach (2006) argue against the postulation of a geometric module on the grounds that the evidence from rats (and infants) can be explained by the fact that subjects only use distant cues for reorientation. By itself this finding does not show that subjects' performance is not best explained by postulating a module. Much depends on whether lost rats cannot use nearby featural cues for reorientation or whether they intelligently choose to ignore such cues as likely to be unreliable. If lost rats cannot use nearby featural cues, we need to know why not and explaining this may involve a module.

<sup>2</sup> Currie and Sterelny, 2000, p. 147; they cite Fodor (1983, p. 71): '[t]he informational encapsulation of the input systems is ... the essence of their modularity.' As already mentioned, Scholl and Leslie (1999b) also hold this view.

properties of modules, there are cases such as Cheng's in which information encapsulation is precisely what modularity is invoked to explain.

Modules are also supposed to play a role in explaining development. For example, three-month-old infants expect unsupported objects to fall and appear to have a relatively sophisticated understanding of how objects can be supported. Infants are also sensitive to other mechanical relations such as pushing and blocking.<sup>3</sup> Postulating a module as the basis for these sensitivities (as, for example, Leslie 1994 does) is supposed to explain how infants, despite being deficient in general knowledge and powers of reasoning, can be sensitive to mechanical relations. It is also supposed to explain why infants' early sensitivity to mechanical relations is limited in some surprising ways,<sup>4</sup> and to play a role in explaining how they later achieve a vast body of knowledge about mechanical relations. Whether this particular claim about modularity is true or false, it illustrates how modules are supposed to explain development.<sup>5</sup> If modularity is to play this explanatory role, it cannot be right to suppose that domain specificity is a feature of modularity just by stipulation because in this case domain specificity is part of what needs explaining.

In addition to cognitive organisation and development, modules are supposed to play a role in explaining patterns of cognitive impairment including double dissociations (Fodor, 1983, pp. 99–100; Shallice, 1988). As we would expect given that the notion of modularity is genuinely explanatory, the existence of a module is usually only one of several candidate explanations for any particular double dissociation (Plaut, 1995). So how are modules supposed to explain patterns of impairment? Let's suppose, just for the sake of illustration, that face recognition is doubly dissociated from object recognition (that is, it's possible for either of these recognitional abilities to be impaired while the other remains intact).<sup>6</sup> How could postulating modules for face and object recognition help to explain this? Modules are often supposed to have a fixed neural architecture; which means, minimally, that the operation of a single module nearly always involve the same neuroanatomically specifiable parts of the brain (e.g. Fodor, 1983, p. 118). It may be tempting to treat this as a stipulation about modularity and to think that this

<sup>3</sup> On infants' understanding of support and blocking, see Baillargeon, Kotovsky and Needham, 1995, pp. 81–3, 95; for pushing see Wang *et al.*, 2003.

<sup>4</sup> See Baillargeon, 2001, pp. 351–2 for a summary of six anomalies.

<sup>5</sup> Is infants' sensitivity to mechanical relations among objects modular? Baillargeon (2002) describes a series of rapid, almost month-by-month, developments in infants' understanding of how objects behave between around 2.5 and 16 months. She argues that better understanding 'is triggered by exposure to unpredicted outcomes' (Lesson 4) and depends 'in part on exposure to appropriate outcome and condition data' (Lesson 5); and that new learning is 'integrated with infants' prior knowledge' (Lesson 6). These findings initially appear inconsistent with a modular account, but Baillargeon endorses at least one account of the mechanisms involved (see Wang, Baillargeon and Paterson 2005, pp. 3–6) which could be used to reconcile modularity with the rapid, experience-driven pattern of development in infancy.

<sup>6</sup> This claim certainly needs careful qualification and may be false; see Bukach *et al.*, 2006 and Kanwisher and Moscovitch, 2000, pp. 3–4 for reviews taking opposing sides.

stipulation makes it possible to explain patterns of cognitive impairment by postulating modules. Where a lesion or developmental disorder results in face or object agnosia, this can be explained by supposing that the neural basis of the corresponding module has been damaged; double dissociations then arise from the possibility of damage to one module while the other is spared.

The role of modules in explaining patterns of impairment can't be nearly this simple. In part this is because modules may involve several neuroanatomically distinct regions of the brain. Face detection, for example, appears to involve at least three distinct regions of the inferior temporal cortex with different but related specialisations (Moscovitch and Moscovitch, 2000, p. 216). More generally, factors such as degeneracy (Price and Friston, 2002) suggest that there is no simple correspondence between modules and neuroanatomical form. More importantly, the simple modular explanation wrongly assumes that the effects of developmental disorders are just like the effects of lesions or strokes. It turns out that developmental disorders affecting face recognition result in patterns of impairment quite different from those caused by lesions to adult brains,<sup>7</sup> and, more generally, that damage to the brain in infancy may result large differences in the functions of anatomically defined units but subtle cognitive deficits.<sup>8</sup> This makes understanding the role of modules in explaining cognitive impairment extremely difficult to unravel. Certainly the stipulation that modules have fixed neuroanatomical forms is insufficient. If modules exist, hypotheses involving modularity must sometimes provide better explanations for patterns of cognitive impairment, both acquired and developmental, than competing nonmodular hypotheses. For this to be possible the notion of modularity must consist in more than a cluster of features. Rather than building stipulations about the neural basis of modules and associated patterns of breakdown into the notion of modularity, we need a notion deep enough to permit investigation of these issues.

To sum up, if modules exist they play a role in explaining cognitive organisation, development and impairment (among other things). But modules could not play these explanatory roles if their possessing information encapsulation, domain specificity or fixed neural architectures were merely a matter of stipulation. So assuming modules exist, there must be more to say about what they are.

---

<sup>7</sup> See de Haan *et al.* (2002, pp. 207–8) who note that 'the infant face-processing system is more widely distributed and/or more plastic following damage than the adult system.' For further contrasts between adult and developmental impairment in face processing, see Johnson, 2005b, pp. 770–2; 2001.

<sup>8</sup> Johnson *et al.* (1998) describe a surprising difference in the effects of brain damage on infants compared to adults, and Karmiloff-Smith *et al.* (2003, pp. 237–8) show that a single cognitive disorder can result in quite different deficits in infancy than in adulthood. Related findings are discussed in (Paterson *et al.* 1999; Frith and Happé, 1998; Tager-Flusberg, 2005). For evidence from computational modelling see Thomas and Karmiloff-Smith (2002): they simulate developmental and acquired disorders by damaging connectionist models of past-tense formation prior to training (for developmental disorders) and after training (for acquired disorders); they find that the two types of damage produce strikingly different outcomes, including differences in what component mechanisms specialise for (2002, pp. 17–18).

## 2. Computation Is the Real Essence of Modularity

I suggest that we can distinguish one kind of cognition from another, and explain the special properties of each kind of cognition, by reference to the kinds of process they involve. Identifying the kind of process modular cognition involves will enable us to explain why modular cognition is domain-specific, informationally encapsulated, and so on. In this section I'll suggest that modular cognition differs from other kinds of cognition in being computational in just the sense in which Fodor once claimed that all cognition is computational. For this to be true, nonmodular forms of cognition such as thinking must not be computational in this special sense; and the hypothesis that modular cognition is computational must explain why modules tend to have features like domain specificity and information encapsulation.<sup>9</sup>

To explain this idea I first need to outline Fodor's general position on thinking and describe an objection to it. Fodor calls his account of thinking the 'Computational Theory of the Mind.' It can be expressed in three words:

*Thinking is Computation* (1998a, p. 9).

What does this mean? Take a computer programmed to detect emotional states expressed in an email message. There are various respects in which this computer could model a human's ability to detect emotions. Here are three, crudely described:

the computer's performance (patterns of successes and failure) matches typical human performance;

the computer's hardware design models biological machinery required by humans for detecting emotional states;

the information and operations of the computer's program resemble a human's states of knowledge and the inferences performed on them.

As I understand it, Fodor's slogan *Thinking is Computation* concerns the third point of comparison. It requires that thinking involves programs whose assertions correspond to what people know and whose operations correspond to inferences people make. To understand how the mind works is to understand how the corresponding programs work. So the claim that thinking is computation amounts to more than the idea that thinking somehow involves computing.<sup>10</sup>

<sup>9</sup> As I understand Fodor the following account of modularity is a minor elaboration of his views, and the key ideas are all drawn from Fodor. That said, I'm not offering its Fodority as a reason for accepting my account and am happy to take credit for anything not Fodor's.

<sup>10</sup> On the Computational Theory of the Mind, see Fodor, 2000, pp. 18–19; Block and Rey, 1998; Fodor, 2005, p. 26 and compare McCarthy, 1998, pp. 22–5; Marr, 1982, pp. 15–30. Note that the Computational Theory invokes a narrow notion of computation. Many models of how the mind works which are appropriately described as computational (e.g. O'Reilly and Munakata, 2000) do not require the Computational Theory of the Mind to be true.

Fodor has recently sketched a potentially sound objection to the Computational Theory of the Mind (Fodor, 2000). If successful, this objection does for the Computational Theory what others have attempted for associative models of thought: it identifies a feature of thought that should but can't be captured by this type of model. As I understand Fodor's argument, one strand of it goes like this:

1. Computational processes are not sensitive to context dependent relations among representations as such.
2. Thinking sometimes involves being sensitive to context dependent relations among representations as such (e.g. the relation ... *is adequate evidence for me to accept that ...*).
3. Therefore, not all thinking is computation.

To follow this argument we need to understand what Fodor means by calling a relation context dependent.<sup>11</sup> In Fodor's terminology, a relation between representations is *context dependent* if whether it holds between two of your representations may depend, in arbitrarily complex ways, on which other mental representations you have. For our purposes, what matters is that the relation ... *is adequate evidence for me to accept that ...* is a context dependent relation. This is because almost anything you know might be relevant to determining what counts as adequate evidence for accepting the truth of a conclusion. Knowing that Sarah missed the conference is (let's suppose) adequate evidence for you to conclude that she is ill ... until you discover that she couldn't resist visiting a cheese factory, or that she urgently needs to finish writing a paper. So the adequate evidence relation is context dependent. But since thinking requires sensitivity to whether evidence is adequate, some of the processes involved in thinking must be sensitive to context dependent relations. So not all of the processes involved in thinking could be computational processes of the kind Fodor envisages. This is why the Computational Theory fails as an account of how we think.

To see how Fodor's argument is supposed to work, compare a similar anti-associative argument about learning:

1. Associative learning processes do not involve retrospective re-evaluation.
2. Learning sometimes involves retrospective re-evaluation.
3. Therefore, not all learning is associative.

To understand this argument we need to know what retrospective re-evaluation is. Suppose you first learn that A-type events and B-type events are both weakly associated with a light coming on. Then, later, you learn that B-type events sometimes happen without the light coming on. As a result of this you decide that the A-type

---

<sup>11</sup> Note that Fodor's use of 'context' does not straightforwardly match any of the uses of 'context' in AI literature as surveyed by Giunchiglia and Bouquet (1997) and it isn't McCarthy's (1986) notion.



events are strongly associated with the light coming on. This counts as retrospective re-evaluation because you're learning about the As when only Bs (and no As) are currently present (Shanks, 2004, p. 233; Dickinson, 2001a, p. 11). It's arguably a feature of associationist models that you can only learn about a cue when it is present, in which case retrospective re-evaluation is incompatible with associationist models.

In both cases, computation and association, the arguments aim to demonstrate that a type of model is inadequate by identifying a feature of thought or learning the model can't capture. Both arguments depend on a complicated mesh of empirical and theoretical issues (on the associative argument, see Dickinson, 2001b; on the computational argument, see McCarthy, 1998). Since establishing this kind of argument requires painstaking theoretical and experimental work rather than philosophical commentary, I won't attempt to convince you that Fodor's argument is correct. But in case Fodor's argument appears obviously wrong or obviously right, let me quickly mention five common responses to it.

First, in Fodor's work this argument is mixed up with claims that current cognitive science is in trouble because the Computational Theory of the Mind is false. Opponents respond that this is irrelevant because the Computational Theory invokes only one of several notions of computation used by cognitive scientists to model the mind (e.g. Pinker, 2005, p. 14). This response, if correct, shows that Fodor's general claims about cognitive science may be false but it is not an objection to the argument presented above, which is only about whether one particular notion of computation can model thinking in one particular way.

Second, some people question whether computational operations are really insensitive to context dependent relations. This is sometimes based on confusion. Consider a semantic theory for a programming language (for an introduction see Finkel, 1996, chapter 10). Such theories need make no reference to context dependent relations; there are no primitive operations like 'discard list items irrelevant to the problem at hand' or 'do this while the evidence is insufficient' which might require context dependent semantics. In Fodor's argument against the Computational Theory, this is all it means to claim that computational operations are insensitive to context dependent relations. It doesn't follow that you can describe why the program exists without reference to context dependent relations. For instance, ... *is relevant to my finding out whether* ... is a context dependent relation between webpages and propositions (since relevance depends what I already know). Given the right phrases, good search software can reliably return results which bear this relation to a specified proposition (as Pinker, 2005, p. 8 emphasises). We need context dependent relations to describe why the software exists but not to give a semantic explanation of how it works. The Computational Theory of the Mind is about the connection between the semantics of programs and the semantics of thought, so it's not relevant whether computers can be useful in gaining knowledge involving context dependent relations (clearly they can), but only whether the semantics of their programs involve such relations (clearly they don't).

Third, instead of objecting that computers are sensitive to context dependent relations, some critics object that humans aren't (e.g. Pinker, 2005, pp. 9–10). This

objection is occasionally based on the following misunderstanding: either humans are infallibly sensitive to context dependent relations or else the appearance of such sensitivity is really just the operation of smart heuristics which do not involve any actual context dependence. The first alternative is clearly untenable. It doesn't follow, however, that our apparent sensitivity to context dependence is just the operation of heuristics. To say that it does follow would be like saying that since we cannot always perceptually discriminate faces from non-faces, what we perceive are not faces but only things with face-like appearances. That we perceive faces imperfectly does not imply that we do not perceive them at all; similarly, that we are imperfectly sensitive to context dependent relations does not imply that we are not sensitive to them at all (see further Fodor, 2005, pp. 29–30).

A fourth response to Fodor's argument is to assert that cognitive scientists have already created computational models of abductive reasoning (where abductive reasoning is the paradigm case requiring sensitivity to context dependent relations). In fact, research in this area divides into two approaches. On the one hand there is research into modelling patterns of reasoning partly for the sake of better understanding reasoning. This approach is exemplified by John McCarthy's work on circumspection (1998, chapters 9, 16), which formalises the principle 'assume you know everything relevant to the problem at hand' and so demonstrates that much is possible without sensitivity to context dependent relations as such. On the other hand, there is research driven by more applied goals such as wanting to integrate textual information with databases or maps (e.g. Dimopoulos and Kakas, 2001). Here algorithms for abduction are sometimes presented, but these algorithms relate to reasoning in humans just as simulated annealing algorithms (Kirkpatrick, Gelatt and Vecchi, 1983; Cerny, 1985) relate to annealing in materials: each is inspired by a phenomenon, not an attempt to capture the theoretically interesting details of it.<sup>12</sup> This research is relevant in more than one way to evaluating Fodor's argument; for example, some of it may support the idea that large parts of human thinking could be accomplished without sensitivity to context dependent relations as such. So far, however, we have no decisive arguments for or against this part of Fodor's argument.

As a final response to Fodor's argument, consider the fact that computers, being able to represent any relation, can represent context dependent relations. Does this straightforwardly undermine Fodor's argument? Compare context dependence with semantic constituency (as in 'dog food' is a *semantic constituent* of 'dinosaur bones found in dog food'). In our ordinary thinking we're sensitive to semantic constituency and this sensitivity arguably couldn't be modelled just by representing facts about constituency: it wouldn't do to list facts about what is a constituent of what. Rather,

---

<sup>12</sup> For example, Dimopoulos and Kakas (1996) describe 'abductive methods' as tools for translating 'high-level specification of goals to low level computational tasks.' Given that they are not aiming to study abductive reasoning, it's no objection to their view that their algorithms need to be told in advance which propositions are candidates for 'best explanation' (see Dimopoulos and Kakas, 1996, §4.2).

we need a system of representation in which there are relations among the vehicles of representation corresponding to relations of semantic constituency among the contents of representation. If Fodor's objection is correct, then it must be similarly impossible to model our sensitivity to evidential relations by representing them: instead, we need a process that is sensitive to context dependent relations between representations. Since it's not in general possible to achieve sensitivity to relations between representations by representing those relations, it isn't obvious that this line of response to Fodor's argument will be successful.

As far as I know, then, there is no simple devastating reply to Fodor's argument against the Computational Theory of the Mind. Equally, the above lines of response suggest that this argument is not obviously correct and depends on a complicated mesh of empirical and theoretical issues. This is not to say we should be agnostic about Fodor's argument. It has considerable force because on the simplest, most straightforward ways of understanding computation and thinking which are not known to be incorrect, the argument's premises come out true. (Similarly, the anti-associative argument mentioned above is forceful not because it's irrefutable but because on the simplest, most straightforward ways of understanding associative mechanisms and learning not known to be false, its premises come out true.) So grant me, if only for the sake of seeing what follows, that Fodor's argument is correct.

In that case the Computational Theory of Mind may succeed as a theory about how the mind copes with syntax or discerns indications of three-dimensional phenomena from retinal stimuli, but it's altogether wrong about what we infer from Sarah's absence from the conference. As Fodor says:

... the Computational Theory is probably true at most of only the mind's modular parts. ... a cognitive science that provides some insight into the part of the mind that isn't modular may well have to be different, root and branch (2000, p. 99).

As this quote implies, Fodor still thinks that some cognitive processes—the modular ones—are computational. I want to suggest that he's right, and also that modular cognition is *essentially* Computational. (From here on I'll write Computational with a capital 'C' as a reminder that our concern is whether the information and operations of a particular kind of computational process can be used as a model for the thoughts and inferences involved in reasoning, not with broader issues about computational models of thinking.)

This proposal departs from Fodor's overall strategy. Fodor starts by asking what thinking is, and answers that it's a special kind of Computational process. He then runs into the awkward problem that such Computation only happens in modules, if at all. Instead of taking this line, we started by asking what modularity is. The answer I'm suggesting is that modular cognition is a Computational process. On this way of looking at things, that such Computation only happens in modules is a useful result because it enables us to identify what is distinctive of modular cognition.

But does being Computational explain modular cognition's characteristic properties? Several forms of explanation are relevant. Perhaps the simplest to describe is David Marr's argument from software design (1982, pp. 102, 325–6). Building Computational systems that are reliable, robust and incrementally improvable involves following certain design principles. These principles are specific to Computational systems and not necessarily relevant to building other kinds of complex system. They require complex tasks to be broken into smaller parts and performed by independent units; the units are to be independent in the sense that one unit's failure will not normally cause other units to fail, communication between units is minimal, and potentially concurrent units do not share state or resources (e.g. Armstrong, 2003, pp. 32–7). These units are archetypal modules and the design principles make precise notions like information encapsulation. Since evolution involves incremental changes, and since evolved systems are required to be reliable and robust no less than designed ones, it is plausible that the same principles will be true of evolved Computational systems. This is one reason for thinking that cognition that is Computational will have properties generally attributed to modules including informational encapsulation.

Another form of explanation linking Computation to properties characteristic of modularity appeals to the fact that Computational processes are not sensitive to context dependent relations. Take domain specificity. If modules cannot process context dependent relations and if notions like evidence and relevance are irreducibly context dependent, then it's perhaps hard to see how modules could be useful. After all, doesn't this mean that the outputs of modular cognition will not be supported by evidence nor relevant to our purposes? Not quite, because when we restrict our attention to a particular domain, evidential and other context-dependent relations become more nearly reducible to relations that are not context dependent. Domain specificity enables Computational processes to be useful despite their insensitivity to context dependent relations. To illustrate, contrast the question, Is this the same *thing* as that? with the question, Is this is the same *face* as that? Without restriction to a domain, it is impossible for a process (or person) insensitive to evidential relations to reliably discern sameness or difference; but some ways of restricting the problem domain make this possible. In general, a Computational process will have to be domain-specific (to some extent) in order to compensate for its insensitivity to context dependent properties.

In addition to restricting a module's operation to a particular domain, approximating evidential and relevance relations with relations that are not context dependent will require restricting the type of input the module is able to process. (Contrast the question, What *in general* counts as evidence that this is the same face as that? with the question, Which *featural information* counts as evidence that this is the same face as that?) This contributes to explaining why a Computational process is likely to be informationally encapsulated (to some extent): insensitivity to context dependent relations limits the range of inputs it can usefully accept.

For a deeper explanatory connection between Computation and information encapsulation, we need to consider the nature of modular representations.

In Section 3, I'll argue that there cannot be direct representational relations between modular cognition and general reasoning. If correct, this argument explains why connections between Computational processes and general reasoning are difficult to establish and are likely to be the exception rather than the norm. This enables us to explain why modules tend to be informationally encapsulated, and—going in the other direction—also why other modules and general reasoning processes have limited access to representations in modules.

In short, there are explanatory connections between Computation and core features of modularity. As you'd expect given that we're talking about a connection between a real essence and superficial features, and given that 'the notion of modularity ought to admit of degrees' (Fodor, 1983, p. 37), these explanations are non-deductive and depend on many further factors. I conjecture that certain systems tend to have various properties characteristic of modules *because* the processes involved are Computational. If so, being a Computational process is what distinguishes modular cognition from other forms of thinking and what explains why modular cognition has properties like domain specificity and information encapsulation. Computation is the real essence of modularity (or, if there are multiple kinds of modularity, Computation is the real essence of one).

Admittedly there are big gaps in my case for this conjecture. Claims about modularity can hardly float free of genes or neurones, but at present I do not know of explanatory connections. Furthermore, recent research in developmental neuroscience suggests that some of the best candidates for modules are not products of maturation and cannot be treated as developmental primitives.<sup>13</sup> Accordingly, a more adequate account of modularity would include a developmental component.

So is Computation really the essence of modularity? I recommend this as a plausible working hypothesis. Not because the case for it is overwhelming or even complete, but because it's as well supported and more open to refutation than any current alternative.

### 3. The Role of Modules in Development

In section 2, I suggested that there is a fundamental way to distinguish types of cognition, which is to identify the kinds of process they involve; and that in the case of modular cognition this is a special kind of computation. Next I want to consider consequences of this approach for the role of modules in development. What follows will depend only on the claim that modular cognition is a different

---

<sup>13</sup> For a review see Johnson, 2005a especially pp. 171–5. For debate over whether and in what sense modules must be innate, see Karmiloff-Smith, 1994 and commentaries, plus her 1998, and Fodor, 1998b, chapters 11, 12. For a case-study, see Karmiloff-Smith, Klima *et al.* 1995; more generally, Johnson, 2000 is a brief discussion of 'bidirectional interactions between brain structure and psychological function' in neural development.

kind of process from ordinary thinking and reasoning, not on the more specific claim that modular cognition is Computational.

Many developmentalists think of modules as 'a basic infrastructure for knowledge and its acquisition' (Wellman and Gelman, 1998, p. 524). Modules are supposed to enable infants and children to learn things that would be too hard to learn just on the basis of experience and general-purpose reasoning. But how do modules facilitate development? This question is 'one of the central issues in developmental psychology' (Carey, 1995, p. 307).

Alan Leslie suggests a simple and apparently attractive answer to these questions: modules facilitate development by providing knowledge. In his words, they 'provide an automatic starting engine for encyclopaedic knowledge' (Leslie, 1988, p. 194). For instance, a module that detects causal relations contributes to development by providing us with knowledge that there are certain causal relations in our environment. This knowledge can then be used for making inferences and guiding action, just as any other knowledge can:

The module ... automatically provides a *conceptual identification* of its input for central thought ... in *exactly the right format* for inferential processes (Leslie, 1988, pp. 193–4 *my italics*).

Here Leslie makes explicit what is nearly always taken for granted, namely that modular cognition results in representations just like full-blown thoughts which are composed of the very concepts we use in general reasoning.<sup>14</sup> On this view, there are direct representational relations between modules and thought.

Elizabeth Spelke has a different account of how modules facilitate development that also requires direct representational relations between modules and thought. Spelke holds that we acquire new abilities 'by *assembling* in new ways the representations delivered by core systems' (Spelke, 2000, p. 1233 *my italics*). For instance, there is evidence that, like Cheng's rats (see Section 1 above), humans have a module that enables them to navigate. Furthermore, younger children are somewhat like rats in being unable to use nearby nongeometric information to navigate when disorientated (Hermer and Spelke, 1996; see also Nadel and Haubach, 2006), whereas adults can combine many kinds of information in locating themselves. Spelke's hypothesis is that the navigation module is essential for development because it's by *assembling* representations from this module with other representations that we become flexible navigators. Elsewhere Spelke also suggests that humans acquire concepts of cardinal numbers by *assembling* representations from an approximate numerosity module with representations

---

<sup>14</sup> Others who accept or tacitly invoke this assumption include Carey (1995, p. 272), Charman and Baron-Cohen (1995, p. 288), Fodor (2000, p. 97), Gerrans (2002, pp. 259–60), Spelke (2003, p. 301), and Sperber (1994, pp. 48–9).

from a core system for representing concrete objects (Spelke, 2000). Her general claim, based on these and further case studies, is that:

The building blocks of all our complex representations are the representations that are constructed from individual core knowledge systems (Spelke, 2003, p. 307).

So where Leslie thinks of modules as providing knowledge, Spelke holds that infants begin life equipped with fragments of knowledge and development consists in *assembling* these fragments into a more unified picture of our environment and its workings.

Put in this way, Spelke's position is schematic until we understand what assembling representations from different core systems involves. Spelke offers at least two explications. One describes assembling representations as bringing together items of specialist knowledge from different domains, as happens for example when scientists discover ways of applying mathematics to describe physical phenomena. On this view, 'conceptual change in childhood is the same sort of process as is conceptual change in the history of science' (Carey and Spelke, 1994, p. 193; cf. Spelke, 1999). A second way to explicate the notion of assembling representations involves language. Here the idea is that language serves as a general-purpose mechanism for combining representations, and assembling ideas means something like forming a phrase out of several words. Spelke illustrates this idea for the case of children discovering how to use both geometric and non-geometric cues in navigation: 'Once they have learnt these terms ["left" and "blue"], the combinatorial machinery of natural language allows children to formulate and understand expressions such as *left of the blue wall* with no further learning' (Spelke, 2003, p. 296).

Whichever way we understand Spelke's notion of assembling, we must assume that the concepts involved in core knowledge are just like the concepts involved in full-blown knowledge, and that these concepts are capable of being combined in a system like language. Spelke is explicit on this point: 'core systems are conceptual and provide a foundation for the growth of knowledge' (Carey and Spelke, 1996, p. 520). So although she takes a very different approach from Leslie, Spelke likewise assumes a direct representational relation between modules and thought.

But suppose it's true, as I claim, that modular cognition and thinking differ in the kinds of process they involve. It would seem to follow that the concepts and representations involved in each case must also differ in kind. Take the concept OBJECT. Here are two schematic notions of what this concept is:

The concept OBJECT is ...

- (a) that in virtue of having which we are able to think about objects as such;
- (b) that in virtue of having which we are able to Compute information about objects as such.

It's standard to suppose that these two notions pick out the same thing. This is because it's standard to suppose that thinking just is Computation. However, this assumption is false if Computation is the real essence of modularity or if modular cognition and thinking are different kinds of process. In that case, it is plausible that the concepts involved in modular representation differ in kind from those involved in ordinary thinking. This is because at least part of what makes concepts the things they are is the fact that they normally feature in certain kinds of process and not others. Where processes differ in kind, so must the concepts. And the same is true for representations.

Contra Leslie and Spelke, then, there are no direct representational relations between modules and thoughts. This greatly complicates how we are to understand the role of modules in development. The complication is not that there cannot be any representational links at all between modular cognition and thinking. Vehicles of representation can be shared for the purpose of transmitting information between systems operating with different kinds of representation. But in such cases, mere sameness of vehicle does not automatically guarantee that information sharing is possible. Indeed, sameness of vehicle guarantees chaos unless there is a mechanism or interface to ensure that the semantic properties each system assigns to shared vehicles are appropriately related. Consequently, if we postulate representational links involving different kinds of representation we must also explain how the representational properties of one system are coordinated with those of the other.

This is not to say that Leslie's or Spelke's approaches are wrong, only that they are missing a crucial step. If some modules are input systems which 'present the world to thought' (Fodor, 1983, p. 101), it may be possible to explain how representations are coordinated in some cases. But how this explanation goes and whether it applies to all modules is an open question.<sup>15</sup>

There is an alternative way to understand the role of modules in development, one that doesn't require modular cognition to provide infants with conceptual identifications of inputs suitable for thinking and reasoning. On this alternative approach, modular cognition assists development by directing and constraining infants' perception and action, thereby enabling them to acquire concepts and knowledge.

An illustration should help to clarify how this idea might be developed. Compare the relation between modular cognition and thinking to the relation between

---

<sup>15</sup> An anonymous reviewer suggested that 'a dichotomous definition of module/non-module may make the problem seem harder than it is.' In fact my account allows for degrees of modularity: a cognitive system may exhibit the features characteristic of modularity to a greater or lesser degree; a Computational model may more or less perfectly fit a cognitive system; and, on at least some accounts of computation (e.g. Copeland, 1996), it is even possible for a process to be Computational to a greater or lesser degree. Nevertheless, it seems to me that there are some cases where cognition involves interactions between very modular and very nonmodular cognitive systems, and these interactions cannot involve direct representational links.



associative learning and practical reasoning. It's clear that the outcomes of associative learning must be to some extent co-ordinated with our practical reasoning otherwise we would be at risk of continually desiring things we are averse to. An approach analogous to Spelke's or Leslie's would require identifying some representational link between the contents of aversions and the contents of desires, so that aversions or their counterparts could feature in practical reasoning. But since the aversions involved in associative learning are unlike desires in being non-propositional, it is extremely difficult to understand how there could be a representational link here—how aversions, which are representations that belong in associative mechanisms, could also feature in practical reasoning. Fortunately Tony Dickinson and Bernard Balleine offer an alternative (Balleine and Dickinson, 1998; Dickinson and Balleine, 2000, p. 193). They suggest that aversions have an effect on our desires only indirectly: they cause physiological changes which we become aware of as affective responses to objects, and this awareness motivates us to modify our desires. Here physiology serves as a non-representational link between the states of aversion and desire and between the processes of conditioning and practical reasoning. In schematic terms, one system has non-representational effects and awareness of these effects feeds into another representational system. The upshot is that the second system operates almost as if it knew what the first system learns, but there are no direct links between representations in the two systems.<sup>16</sup>

Dickinson and Balleine's idea about aversion's relation to desire can be extended to modular cognition's relation to thought. Take concept acquisition. A necessary condition for acquiring many concepts is the ability to identify similarities among the things that fall under them. Suppose we have a module capable of detecting some such similarities. One way for this module to facilitate concept-acquisition is by causing us to respond in similar ways when presented with objects which are similar in the relevant respects. In such cases we will have two kinds of similarity in play: similarities among the objects and similarities among our reactions to them.<sup>17</sup> The similarities in our reactions to the objects could draw our attention to the similarities in the objects themselves, thereby assisting us in acquiring the concept. This is one way in which a module could help us to acquire concepts by virtue of directing and constraining our perception, action or attention.

In general, whereas Leslie's and Spelke's approaches require representational connections between modules and thought, I suppose that modules may facilitate development via non-representational connections. The difference between their approaches and mine is very loosely analogous to the difference between conveying

<sup>16</sup> For evidence that associative mechanisms influence humans' judgements, see Dickinson, 2001a and Shanks (see e.g. Shanks, 2004, p. 234).

<sup>17</sup> The importance of such pairs of similarities for concept learning has been stressed by Donald Davidson (e.g. Davidson, 1991, p. 200).

information verbally (i.e. providing a representation) and conveying information by nonverbally pointing someone to a source of information.<sup>18</sup>

You might object that the envisaged role for modules would be ineffectual given the tremendous difficulty infants and children face in acquiring concepts and learning about the world. This is a fair objection to the bare idea that modules cause us to react similarly to things that are similar in respects relevant to conceptual classification. But we can elaborate this idea by appealing to evidence that modular cognition may result in eye movements (Clements and Perner, 1994; Clements, Rustin and McCallum, 2000; Carey and Spelke, 1996, p. 522) and direct our attention (Leslie, Xu *et al.*, 1998; Scholl and Leslie, 1999a; Carey and Xu, 2001).<sup>19</sup> If modules can move our eyes and direct our attention, it's plausible that they can play a rich and essential role in development without representational links between modular cognition and thought.

One relatively clear case where modules appear to play an indirect role in conceptual development is speech. Infants enjoy categorical perception of phonemes from four months or earlier (Eimas, Siqueland *et al.*, 1971), which arguably involves a speech module (Liberman and Mattingly, 1985). By contrast, the acquisition of phoneme concepts as measured by standard tests for phonological awareness (e.g. Anthony and Lonigan, 2004, p. 46) takes several years, varies systematically depending on oral language, and is differentially facilitated by different writing systems (Anthony and Francis, 2005, pp. 256, 257). Furthermore, certain distinctions between phonemes are hard to identify conceptually despite being unproblematically perceived (Treiman, Broderick *et al.*, 1998). Apparently, then, infants' modular cognition of speech does not facilitate development by providing conceptual identifications and plays a significant but indirect role in facilitating the development of capacities to think and reason about speech.

This section has been very tentative. Here's what's not tentative: (1) several views about the role of modules in development assume that there are direct representational relations between modules and thought; (2) this assumption is implausible if modular cognition and thinking are different kinds of process (for instance because modular cognition is unlike thought in being Computational);

---

<sup>18</sup> After writing this paper I found that Gelman and Leslie (2001, p. 61) endorse but do not develop a possibly similar position: 'Specialized cognitive mechanisms can allow attention to go to specific properties or sets of properties. ... For example, the "Michotte module" allows young infants to attend to causes and effects, grounding the concept "cause" without infants knowing anything about what causes really are.'

<sup>19</sup> These studies show that the control of attention depends on processing information about features that are not conceptually identified in thought; in the case of infants, the subject may even be incapable of thinking about these features. I take this to be part of the case for the claim that modules may have a non-representational yet rich role in facilitating development. Some of the authors make the opposite claim. According to them, the module or system involved in controlling attention 'creates the representations that articulate thought. That is, it preempts ... representations in our experiences of the world' (Carey and Xu, 2001, p. 203).

and (3) if this assumption is false, we need to re-examine how modules facilitate development.

#### 4. Conclusion

I have suggested that in order to understand what modules are we need not only to describe the cluster of properties characteristic of modularity but also to explain why modules have these properties. We can do this by identifying the kind of process distinctive of modular cognition, and this process is a special kind of computation.

If modular cognition and thinking differ in the kinds of process they involve, there cannot be direct representational links between modules and thoughts which explain how modules facilitate development. Instead, at least part of the role of modules in explaining development may involve guidance by eye movements, behaviour and attention. On this view, modules function like nonverbal instructors, pointing us to objects and enabling us to track similarities.

As stressed earlier, the claim that modular cognition is a special kind of computation might easily be wrong. I hope at least to have explained why the question about what modules are is pressing and what an answer to it should achieve. If modules play a central explanatory role in theories about how infants and children acquire knowledge of people and things, we need a deeper understanding of what modular cognition is and how it differs from, and interacts with, other forms of cognition.

*Department of Philosophy  
University of Warwick*

#### References

- Anthony, J.L. and Francis, D.J. 2005: Development of phonological awareness. *Current Directions in Psychological Science*, 14(5), 255–260.
- Anthony, J.L. and Lonigan, C.J. 2004: The nature of phonological awareness: converging evidence from four studies of preschool and early grade school children. *Journal of Educational Psychology*, 96(1), 43–55.
- Armstrong, J. (2003), *Making Reliable Distributed Systems in the Presence of Software Errors*. Stockholm: Royal Institute of Technology.
- Baillargeon, R. 2001: Infants' physical knowledge: of acquired expectations and core principles. In E. Dupoux (ed.), *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, MA: MIT Press.
- Baillargeon, R. 2002: The acquisition of physical knowledge in infancy: a summary in eight lessons. In U. Goswami (ed.), *Blackwell Handbook of Childhood Cognitive Development*. Oxford: Blackwell.

- Baillargeon, R., Kotovsky, L. and Needham, A. 1995: The acquisition of physical knowledge in infancy. In D. Sperber and D. Premack (eds), *Causal Cognition. A Multidisciplinary Debate*. Oxford: Clarendon.
- Balleine, B. and Dickinson, A. 1998: Consciousness—the interface between affect and cognition. In J. Cornwell (ed.), *Consciousness and Human Identity*. Oxford: Oxford University Press.
- Baron-Cohen, S. 1995: *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA; London: MIT Press.
- Block, N. and Rey, G. 1998: Mind, computational theories of. In E. Craig (ed.), *Routledge Encyclopedia of Philosophy*. London: Routledge.
- Bukach, C.M., Gauthier, I. and Tarr, M.J. 2006: Beyond faces and modularity: the power of an expertise framework. *Trends in Cognitive Sciences*, 10(4), 159–166.
- Buller, D.J. and Hardcastle, V.G. 2000: Evolutionary psychology, meet developmental neurobiology: against promiscuous modularity. *Brain and Mind*, 1, 307–325.
- Carey, S. 1995: On the origin of causal understanding. In D. Sperber, D. Premack and A.J. Premack (eds), *Causal Cognition: A Multidisciplinary Debate*. Oxford: Clarendon.
- Carey, S. and Spelke, E. 1994: Domain-specific knowledge and conceptual change. In L. Hirschfeld and S. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge: Cambridge University Press.
- Carey, S. and Spelke, E. 1996: Science and core knowledge. *Philosophy of Science*, 63, 515–533.
- Carey, S. and Xu, F. 2001: Infants' knowledge of objects: beyond object files and object tracking. *Cognition*, 80, 179–213.
- Carruthers, P. 2006: The case for massively modular models of mind. In R.J. Stainton (ed.), *Contemporary Debates in Cognitive Science*. Oxford: Blackwell.
- Cerny, V. 1985: Thermodynamical approach to the traveling salesman problem: an efficient simulation Algorithm. *Journal of Optimization Theory and Applications*, 45, 41–51.
- Charman, T. and Baron-Cohen, S. 1995: Understanding photos, models and beliefs: a test of the modularity thesis of theory of mind. *Cognitive Development*, 10, 287–298.
- Cheng, K. 1986: A purely geometric module in the rat's spatial representation. *Cognition*, 23, 149–178.
- Clements, W. and Perner, J. 1994: Implicit understanding of belief. *Cognitive Development*, 9, 377–395.
- Clements, W., Rustin, C. and McCallum, S. 2000: Promoting the transition from implicit to explicit understanding: a training study of false belief. *Developmental Science*, 3(1), 81–92.
- Coltheart, M. 1999: Modularity and cognition. *Trends in Cognitive Sciences*, 3(3), 115–120.
- Copeland, B.J. 1996: What is computation? *Synthese*, 108, 335–359.
- Currie, G. and Sterenly, K. 2000: How to think about the modularity of mind-reading. *Philosophical Quarterly*, 50(199), 145–160.

- Davidson, D. 1991: Epistemology externalized. In *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press. Originally published in *dialectica*, 45(2-3): 191–202.
- de Haan, M., Humphreys, K. and Johnson, M.H. 2002: Developing a brain specialized for face perception: a converging methods approach. *Developmental Psychobiology*, 40, 200–212.
- Dickinson, A. 2001a: Causal learning: an associative analysis. *The Quarterly Journal of Experimental Psychology*, 54B(1), 3–25.
- Dickinson, A. 2001b: Causal learning: association versus computation. *Current Directions in Psychological Science*, 10(4), 127–132.
- Dickinson, A. and Balleine, B. 2000: Causal cognition and goal-directed action. In C. Heyes and L. Huber (eds), *The Evolution of Cognition*. Cambridge, MA: MIT Press.
- Dimopoulos, Y. and Kakas, A. 1996: Abduction and learning. In L. De Raedt (ed.), *Advances in Inductive Logic Programming*. Amsterdam: IOS Press.
- Dimopoulos, Y. and Kakas, A. 2001: Information integration and computational logic. *Computational Logic*, Special Issue: Technological Roadmap for CL, 105–135.
- Eimas, P.D., Siqueland, E.R., Jusczyk, P. and Vigorito, J. 1971: Speech perception in infants. *Science*, 171(3968), 303–306.
- Finkel, R. 1996: *Advanced Programming Language Design*. New York: Addison-Wesley.
- Fodor, J. 1983: *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA; London: MIT Press.
- Fodor, J. 1998a: *Concepts*. Oxford: Clarendon.
- Fodor, J. 1998b: *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA; London: MIT Press.
- Fodor, J. 2000: *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.
- Fodor, J. 2005: Reply to Steven Pinker, 'So How Does the Mind Work?' *Mind & Language*, 20(1), 25–32.
- Frith, U. and Happé, F. 1998: Why specific developmental disorders are not specific: on-line and developmental effects in autism and dyslexia. *Developmental Science*, 1(2), 267–272.
- Gelman, T.P. and Leslie, A. 2001: Children's inferences from 'knowing' to 'pretending' and 'believing'. *British Journal of Developmental Psychology*, 19, 59–83.
- Gerrans, P. 2002: Modularity reconsidered. *Language and Communication*, 22(3), 259–268.
- Giunchiglia, F. and Bouquet, P. 1997: Introduction to contextual reasoning. In B. Kokinov (ed.), *Perspectives on Cognitive Science Volume 3*. Sofia, Bulgaria: NBU Press.
- Hermer, L. and Spelke, E. 1996: Modularity and development: the case of spatial reorientation. *Cognition*, 61, 195–232.
- Hirschfeld, L. and Gelman, S. (eds.) 1994. *Mapping the Mind: Domain Specificity in Cognition And Culture*. Cambridge: Cambridge University Press.

- Johnson, M.H. 2000: Functional brain development in infants: elements of an interactive specialization framework. *Child Development*, 71(1), 75–81.
- Johnson, M.H. 2001: Functional brain development in humans. *Nature Neuroscience*, 2, 475–483.
- Johnson, M.H. 2005a: *Developmental Cognitive Neuroscience*, 2nd edn. Oxford: Blackwell.
- Johnson, M.H. 2005b: Subcortical face processing. *Nature Neuroscience*, 6, 766–74.
- Johnson, M.H., Tucker, L.A., Stiles, J. and Trauner, D. 1998: Visual attention in infants with perinatal brain damage: evidence of the importance of anterior lesions. *Developmental Science*, 1(1), 53–58.
- Kanwisher, N. and Moscovitch, M. 2000: The cognitive neuroscience of face processing: an introduction. *Cognitive Neuropsychology*, 17(1–3), 1–13.
- Karmiloff-Smith, A. 1994: Precis of beyond modularity: A developmental perspective on cognitive science. *Behavioral and Brain Sciences*, 17(4), 693–745.
- Karmiloff-Smith, A. 1998: Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences*, 2(10), 389–98.
- Karmiloff-Smith, A. 2006: Modules, genes and evolution: what have we learnt from atypical development? In Y. Munakata and M.H. Johnson (eds), *Processes of Change in Brain and Cognitive Development: Attention and Performance XXI*. Oxford: Oxford University Press.
- Karmiloff-Smith, A., Brown, J.H., Grice, S. and Paterson S. 2003: Dethroning the myth: cognitive dissociations and innate modularity in Williams Syndrome. *Developmental Neuropsychology*, 23(1–2), 227–242.
- Karmiloff-Smith, A., Klima, E., Bellugi, U., Grant, J. and Baron-Cohen, S. 1995: Is there a social module? Language, face processing, and theory of mind in individuals with Williams Syndrome. *Journal of Cognitive Neuroscience*, 7(2), 196–208.
- Kirkpatrick, S., Gelatt, C.D. and Vecchi, M.P. 1983: Optimization by simulated annealing. *Science*, 4598(13), 671–680.
- Leslie, A. 1988: The necessity of illusion: Perception and thought in infancy. In L. Weiskrantz (ed.), *Thought Without Language*. Oxford: Clarendon.
- Leslie, A. 1994: ToMM, ToBY, and agency: core architecture and domain specificity. In L. Hirschfeld and S. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge: Cambridge University Press.
- Leslie, A., Xu, F., Tremoulet, P.D. and Scholl, B. J. 1998: Indexing and the object concept: developing ‘what’ and ‘where’ systems. *Trends in Cognitive Sciences*, 2(1).
- Lieberman, A.M. and Mattingly, I.G. 1985: The motor theory of speech perception revised. *Cognition*, 21(1), 1–36.
- Marr, D. 1982: *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.
- McCarthy, J. 1986: Notes on formalizing contexts. In T. Kehler and S. Rosenschein (eds), *Proceedings of the Fifth National Conference on Artificial Intelligence*. Los Altos, California: Morgan Kaufmann.
- McCarthy, J. 1998: *Formalizing Common Sense: Papers by John McCarthy*. Exeter, UK: Intellect.

- Moscovitch, M. and Moscovitch, D.A. 2000: Super face-inversion effects for isolated internal or external features, and for fractured faces. *Cognitive Neuropsychology*, 17(1-3), 201-219.
- Nadel, L. and Haupbach, A. 2006: Species comparisons in development: the case of the geometric 'module'. In Y. Munakata and M.H. Johnson (eds), *Processes of Change in Brain and Cognitive Development: Attention and Performance XXI*. Oxford: Oxford University Press.
- O'Reilly, R.C. and Munakata, Y. 2000: *Computational Explorations in Cognitive Neuroscience*. Cambridge, MA: MIT Press.
- Paterson, S., Brown, J.H., Gsödl, M.K., Johnson, M.H. and Karmiloff-Smith, A. 1999: Cognitive modularity and genetic disorders. *Science*, 286, 2355-9.
- Pinker, S. 2005: So how does the mind work? *Mind & Language*, 20(1), 1-24.
- Plaut, D.C. 1995: Double dissociation without modularity: evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, 17(2), 291-321.
- Price, C.J. and Friston, K.J. 2002: Degeneracy and cognitive anatomy. *Trends in Cognitive Sciences*, 6(10), 416-22.
- Samuels, R. 1998: Evolutionary psychology and the massive modularity hypothesis. *The British Journal for the Philosophy of Science*, 49(4), 575-602.
- Samuels, R. 2004: Innateness in cognitive science. *Trends in Cognitive Sciences*, 8(3), 136-41.
- Scholl, B.J. and Leslie, A. 1999a: Explaining the infant's object concept: beyond the perception/cognition dichotomy. In E. LePore and Z. Pylyshyn (eds), *What Is Cognitive Science?* Oxford: Blackwell.
- Scholl, B.J. and Leslie, A. 1999b: Modularity, development and 'theory of mind'. *Mind & Language*, 14(1), 131-153.
- Shallice, T. 1988: *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Shanks, D. 2004: Judging covariation and causation. In D. Koehler and N. Harvey (eds), *Blackwell Handbook of Judgment and Decision Making*. Oxford: Blackwell.
- Spelke, E. 1988: Where perceiving ends and thinking begins: the apprehension of objects in infancy. In A. Yonas (ed.), *Perceptual Development in Early Infancy*. Hillsdale, NJ: Erlbaum.
- Spelke, E. 1999: Unity and diversity in knowledge. In E. Winograd, R. Fivush and W. Hirst (eds), *Ecological Approaches to Cognition: Essays in Honor of Ulric Neisser*. Mahwah, NJ: Erlbaum.
- Spelke, E. 2000: Core knowledge. *American Psychologist*, 55, 1233-1243.
- Spelke, E. 2003: What makes us smart? In D. Gentner and S. Goldin-Meadow (eds), *Advances in the Study of Language and Thought*. Cambridge, MA: MIT Press.
- Sperber, D. 1994: The modularity of thought and the epidemiology of representations. In L. Hirschfeld and S. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge: Cambridge University Press.
- Sperber, D. 2001: In defense of massive modularity. In E. Dupoux (ed.), *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, MA: MIT Press.

- Sperber, D. and Wilson, D. 2002: Pragmatics, modularity and mind-reading. *Mind & Language*, 17(1-2), 3-23.
- Tager-Flusberg, H. 2005: What neurodevelopmental disorders can reveal about cognitive architecture: the example of theory of mind. In P. Carruthers, S. Laurence and S. Stich (eds), *The Innate Mind*. Oxford: Oxford University Press.
- Tanaka, J. and Gauthier, I. 1997: Expertise in object and face recognition. In R.L. Goldstone, P.G. Schyns and D.L. Medin (eds), *Mechanisms of Perceptual Learning vol. 36*. San Diego: Academic Press.
- Thomas, M.S.C. and Karmiloff-Smith, A. 2002: Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioral and Brain Sciences*, 25, 1-60.
- Treiman, R., Broderick, V., Tincoff, R. and Rodriguez, K. 1998: Children's phonological awareness: confusions between phonemes that differ only in voicing. *Journal of Experimental Child Psychology*, 68(1), 3-21.
- Wang, S.-H., Baillargeon, R. and Paterson, S. 2005: Detecting continuity violations in infancy: a new account and new evidence from covering and tube events. *Cognition*, 95(2), 129-173.
- Wang, S.-H., Kaufman, L. and Baillargeon R. 2003: Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events. *Infant Behavior and Development*, 26, 529-567.
- Wellman, H. and Gelman, S. 1998: Knowledge acquisition in foundational domains. In D. Kuhn and R. S. Siegler (eds), *Handbook of Child Psychology*. New York: Wiley.