

Is the mind really modular?

Jesse J. Prinz

When Fodor titled his (1983) book the *Modularity of Mind*, he overstated his position. His actual view is that the mind divides into systems some of which are modular and others of which are not. The book would have been more aptly, if less provocatively, called *The Modularity of Low-Level Peripheral Systems*. High-level perception and cognitive systems are non-modular on Fodor's theory. In recent years, modularity has found more zealous defenders, who claim that the entire mind divides into highly specialized modules. This view has been especially popular among Evolutionary Psychologists. They claim that the mind is massively modular (Cosmides and Tooby, 1994; Sperber, 1994; Pinker, 1997; see also Samuels, 1998). Like a Swiss Army Knife, the mind is an assembly of specialized tools, each of which has been designed for some particular purpose. My goal here is to raise doubts about both peripheral modularity and massive modularity. To do that, I will rely on the criteria for modularity laid out by Fodor (1983). I will argue that neither input systems, nor central systems are modular on any of these criteria.

Some defenders of modularity have dropped parts of Fodor's definition and defined modularity with reference to a more restricted list of features. Carruthers (this volume) makes such a move. My arguments against modularity threaten these accounts as well. My claim is not just that Fodor's criteria are not jointly satisfied by subsystems within the mind, but they are rarely satisfied individually. When we draw boundaries around subsystems that satisfy any one of Fodor's criteria for modularity, we find, at best, scattered islands of modularity. If modules exist, they are few and far between. The kinds of systems that have been labeled modular by defenders of both peripheral and massive modularity probably don't qualify. Thus, modularity is not a very useful construct in doing mental cartography.

1. Fodor's Criteria

Modularity should be contrasted with the uncontroversial assumption of "functional decomposition": the mind contains systems that can be distinguished by the functions they carry out. The modularity hypothesis is a claim about what some of the systems underlying human competences are like. Fodor characterizes modularity by appeal to nine special properties. He says that modular systems are:

- (1) Localized: modules are realized in dedicated neural architecture
- (2) Subject to characteristic breakdowns: modules can be selectively impaired
- (3) Mandatory: modules operate in an automatic way
- (4) Fast: modules generate outputs quickly

- (5) Shallow: modules have relatively simple outputs (e.g., not judgments)
- (6) Ontogenetically determined: modules develop in a characteristic pace and sequence
- (7) Domain specific: modules cope with a restricted class of inputs
- (8) Inaccessible: higher levels of processing have limited access to the representations within a module
- (9) Informationally encapsulated: modules cannot be guided by information at higher levels of processing

Fodor's criteria can be interpreted in different ways. Perhaps a system is modular to the extent that it exhibits properties on the list. Alternatively, some of the properties may be essential, while others are merely diagnostic. In recent writings, Fodor (2000) has treated informational encapsulation as a *sine qua non* for modularity. Defenders of massive modality focus on domain specificity in ontogenetic determination (Cosmides & Tooby, 1994; Sperber, 1994). I will emphasize these properties in what follows, but I will also discuss the other properties on the list, because, even if they are not essential, Fodor implies that they cluster together. I am skeptical. I think the properties on Fodor's list neither can be used neither jointly nor individually to circumscribe an interesting class of systems.

2. Localization and Characteristic Breakdowns

The first two items in Fodor's account of modularity—localization and characteristic breakdowns—are closely related. The claim that mental faculties are localized is supported by the fact that focal brain lesions cause selective mental deficits. Further evidence for localization comes from neuroimaging studies, which purport to pinpoint the brain areas that are active when healthy individuals perform mental tasks.

The evidence for anatomical localization seems overwhelming at first, but problems appear on closer analysis. Uttal (2001) points out that there is considerable inconsistency across laboratories and studies. For example, there is little agreement about the precise location of Broca's area, the alleged center of language production (Poeppel, 1996). Indeed, aspects of language production have been located in every lobe of the brain (Pulvermüller, 1999). Or consider vision. There is considerable debate about the location of systems involved in processing things as fundamental as space and color. Uttal also points out that neuroimaging studies often implicate large-scale networks, rather than small regions, suggesting that vast expanses of cortex contribute to many fundamental tasks. Sometimes the size of these networks is underestimated. By focusing on hotspots, researchers often overlook regions of the brain that are moderately active during task performance.

Lesion studies are mired by similar problems. Well-known deficits, such as visual neglect, are associated with lesions in entirely different parts of the brain (e.g., frontal eye-fields and inferior parietal cortex). Sometimes, lesions in the same area have different effects in different people, and all too often neuropsychologists draw general conclusions from individual case studies. This assumes localization rather than providing evidence for it. Connectionist models have been used to show that focal lesions can lead to specific deficits even when there is no localization of functions: a massively distributed artificial neural network can exhibit a selective deficit after a few nodes are removed (simulating a focal lesion), even though those nodes were not the locus of the capacity that is lost (Plaut, 1995). More generally, when a lesion leads to an impairment of a capacity, we do not know if the locus of the lesion is the neural correlate of the capacity or the correlate of some ancillary prerequisite for the capacity.

I do not want to exaggerate the implications of these considerations. There is probably a fair degree of localization in the brain. No one is tempted to defend Lashley's (1950) equipotentiality hypothesis, according to which the brain is an undifferentiated mass. But the rejection of equipotentiality does not support modularity. Defenders of modularity combine localization with domain specificity: they assume that brain regions are exclusively dedicated to specific functions. Call this "strong localization." If, in reality, mental functions are located in large-scale overlapping networks, then it would be misleading to talk about anatomical regions as modules.

Evidence for strong localization is difficult to come by. Similar brain areas are active during multiple tasks, and focal brain lesions tend to produce multiple deficits. For example, aphasia patients regularly have impairments unrelated to language (Bates, 1994; Bates et al., 2001). Even genetic language disorders (specific language impairments) are co-morbid with nonlinguistic problems, such as impairments in rapid auditory processing or orofacial control (Tallal et al., 1996; Vargha-Khadem et al., 1995; Bishop, 1992).

To take another example, consider Stone et al.'s (2002) discussion of a patient who is said to have a selective deficit in reasoning about social exchanges. This patient is also impaired in recognizing *faux pas* and mental state terms, so he does not support the existence of a social exchange module. Nor does this patient support the existence of a general social cognition module, because he performs well other social tasks.

In sum it is difficult to find cases where specific brain regions have truly specific functions. One could escape the localization criterion by defining modules as motley assortments of abilities (e.g., syntax plus orofacial control; social exchange plus *faux pas*), but this would trivialize the modularity hypothesis. There is little evidence that the capacities presumed to be modular by defenders of the modularity hypothesis are strongly localized.

3. Mandatory, Fast, and Shallow

The next three items on Fodor's characterization of modules are supposed to capture a distinctive style of processing. Modules, he says, are mandatory, fast, and shallow. I don't think that these properties capture an interesting class of systems within the mind. There is little reason to think they are intimately related to each other. A system whose processes are mandatory (i.e., automatic) need not be fast. For example, consider the

system underlying circadian rhythms, which regulate the sleep-wake cycle. Nor should we expect mandatory processes to be shallow. Semantic priming is mandatory, but it taps into conceptual knowledge. The three properties under consideration are more of a grab bag than a coherent constellation.

The three properties are uninteresting when considered in isolation. Consider automaticity. Everyone agrees that some mental processes are automatic, but most mental capacities seem to integrate automatic processes with processes that are controlled. For example, we form syntactic trees automatically, but sentence production can be controlled by deliberation. Likewise, we see colors automatically, but we can visually imagine colors at will. The automatic/controlled distinction cannot be used to distinguish systems in an interesting way.

Now consider speed. As remarked above, some capacities that look like plausible candidates for mental modules may be slow (e.g., those governing circadian rhythms). In addition, there are large variations in performance speed within any general system, such as vision or language. Verb conjugation, for example, may depend on whether the verb in question is regular or irregular, and whether the verb is frequent or infrequent. There is little inclination to say that verb conjugation is more modular when it is accomplished more quickly. In addition, some of the worst candidates for modular processes are relatively fast: priming is instantaneous but it can link elements in entirely different systems (the smell of coffee may evoke memories of a holiday in Rome).

Finally, consider the suggestion that modules have shallow outputs. Shallow outputs are outputs that do not require a lot of processing. As an example, Fodor suggests that it doesn't take the visual system much processing to output representations of basic categories (e.g., apple, chair, car). But how much processing is too much? There is a lot of processing between retinal stimulation and visual recognition, and far fewer steps in certain higher cognitive processes, which Fodor regards as nonmodular (e.g., it takes one step to infer "fiscal conservative" from "Republican"). To get around this difficulty, one might restrict "shallow outputs" to nonconceptual outputs. Carruthers (this volume) rightly complains that this would beg the question against defenders of massive modularity: they claim conceptual tasks are modular. Definitions of "shallowness" are either too inclusive or too exclusive. It is not a useful construct for dividing up the mind.

4. Ontogenetic Determinism

Fodor implies that modules are ontogenetically determined: they develop in a predictable way in all healthy individuals. Modules emerge through the maturation, rather than learning and experience. In a word, they are innate. I am skeptical. I think many alleged modular systems are learned, at least in part.

Of all alleged modules, the senses have the best claim to being innate, but they actually depend essentially on experience. Within the neocortex of infants, there is considerably less differentiation between the senses than there is in adults. Cortical pathways seem to emerge through a course of environmentally stimulated strengthening of connections and pruning. One possibility is that low-level sensory mechanisms are innate (including sense organs, subcortical sensory hubs, and the cytoarchitecture of primary sensory cortices), while high-level sensory mechanisms are acquired through

environmental interaction (Quartz and Sejnowski, 1997). This conjecture is supported by the plasticity of the senses (Chen et al. 2002). For example, amputees experience phantom limbs because unused limb-detectors get rewired to neighboring cells, and blind people use brain areas associated with vision to read Braille. In such cases, sensory wiring seems to be input driven. Thus, it is impossible to classify the senses as strictly innate or acquired.

Strong nativist claims are even harder to defend when we go outside the senses. Consider folk physics: our core knowledge of how medium-sized physical objects behave. It is sometimes suggested that folk physics is an innate module. For example, some developmental psychologists conjecture that infants innately recognize that objects move as bounded wholes, that objects, cannot pass through each other, and that objects fall when dropped. I don't find these conjectures plausible (see Prinz, 2002). Newborns are not surprised by violations of boundedness (Slater et al., 1990), and five-month-olds are not surprised by violations of solidity and gravity (Needham & Baillargeon, 1993). Indeed, some tasks involving gravity and solidity even stump two-year-olds (Hood et al., 2000). My guess is that innate capacities to track movement through space combine with experience to derive the basic principles of folk physics (compare Scholl & Leslie, 1999). If so, folk physics is a learned byproduct of general tracking mechanisms.

Consider another example: massive modularists claim that we have an innate capacity for "mindreading," i.e., attributing mental states (e.g., Leslie, 1994; Baron-Cohen, 1995). The innateness claim is supported by two facts: mindreading emerges on a fixed schedule, and it is impaired in autism, which is a genetic disorder. Consider these in turn. The evidence for a fixed schedule comes from studies of healthy Western children. Western children generally master mindreading skills between the third and fourth birthdays in normally developing children. However, this pattern fails to hold up cross-culturally (Lillard, 1998; Vinden, 1999). For example, Quechua speakers of Peru don't master belief attribution until they are eight (Vinden, 1996). Moreover, individual differences in belief attribution are highly correlated with language skills and exposure to social interaction (Garfield et al., 2001). This suggests that mindreading skills are acquired through social experience and language training.

What about autism? I don't think that mindreading deficit in autism is evidence for innateness. An alternative hypothesis is that mindreading depends on a more general capacity which is compromised in autism. One suggestion is that autists' difficulty with mindreading is a consequence of genetic abnormality in oxytocin transmission, which prevents them from forming social attachments, and thereby undermines learned social skills (Insel et al., 1999).

As a final example, I want to consider language. I will keep my remarks brief, because I have criticized the evidence for an innate language faculty elsewhere (Prinz, 2002; see also Pullum and Scholz (this volume). I restrict myself to a brief comment on the allegation that language emerges on a fixed schedule. It seems, for example, that children reliably begin to learn words between eight and ten months, and they begin to combine words around eighteen months. These numbers do not support innateness. They are statistical averages that belie enormous variation. Bates et al. (1995) found that, in early word comprehension, age accounted for only 36% of the variance, and individual differences were huge. In a sample of ten-month-olds, the reported number of words known ranged from 0 to 144. Among eighteen-month-olds, Bates et al. found that 46%

combined words sometimes, and 11% did so frequently. The rate of learning may depend on factors such as general cognitive development and working memory span (e.g., Seung and Chapman, 2000). If the rate of language acquisition is variable and correlated with nonlinguistic factors, then it is bad evidence for innateness.

In presenting these examples, I have been trying to show that the evidence for innateness has been exaggerated. The developmental trajectory of many mental capacities is consistent with a learning story. I don't mean to suggest that we lack specialized capacities. Specialized capacities can be learned. This has been demonstrated by recent work on computational modeling (Jacobs, 1999). For example, one class of connectionist models works on the principle that inputs will be processed in the portion of the network that makes fewest errors when processing the training data. Using such a model, Thomas and Karmiloff-Smith (2002) demonstrate that a network trained to form past-tense verbs from their present-tense forms will spontaneously produce a subcomponent that handles regular verbs and another subcomponent that handles irregulars. Their network has one pathway with three layers of units and another pathway with two. The two-layer pathway is better with regulars and the three-layer pathway is better with irregulars, because irregular endings are not linearly separable. These two pathways are not task specific before training, but they end up being task specific afterwards. Such toy examples show that we can easily acquire specialized subsystems through learning. That means there is no reason to expect an intimate link between innateness and specialization. Once that link is broken, the role of innateness in defending modularity is cast into doubt.

5. Domain Specificity

Domain specificity is closely related to innateness. To say that a capacity is innate is to say that we are biologically prepared that specific capacity. Innate entails domain specific. But, as we have just seen, domain specific does not entail innate. Therefore, in arguing against the innateness criterion of modularity, I have not undermined the domain specificity criterion. Domain specificity is regarded by some as the essence of modularity, and it deserves careful consideration.

It is difficult to assess the claim that some mental systems are domain specific without clarifying definitions. What exactly is a "domain"? What is "specificity"? On some interpretations, domain specificity is a trivial property. "Domain" can be interpreted as a synonym for "subject matter." To say that a cognitive system concerns a domain, on this reading, is to say that the system has a subject matter. The subject matter might be a class of objects in the world, a class of related behaviors, a skill, or any other coherent category. On the weak reading, just about anything can qualify as a domain. Consider an individual concept, such as the concept CAMEL. A mental representation used to categorize camels is specific to a domain, since camels are a coherent subject matter. Likewise for every concept.

"Specificity" also has a weak reading. In saying that a mental resource is domain specific, we may be saying no more than that it is used to process information underlying our aptitude for that domain. In other words, domain specificity would not require exclusivity. Consider the capacity to throw crumpled paper into a wastebasket.

Presumably, the mental resources underlying that ability overlap with resources used in throwing basketballs in hoops or throwing tin cans into recycle bins. On the weak definition of “specificity,” we have a domain specific capacity for throwing paper into wastebaskets simply in virtue of having mental resources underlying that capacity, regardless of the fact that those resources are not dedicated exclusively to that capacity.

Clearly defenders of domain specificity want something more. On a stronger reading, “domain” refers, not to any subject matter, but to matters that are relatively encompassing. Camels are too specific. The class of animals might qualify as a domain, because it is more inclusive. Psychologists have this kind of category in mind when they talk about “basic ontological domains.” But notice that the stronger definition is hopelessly vague. What does it mean to say domains are relatively encompassing? Relative to what? CAMEL is an encompassing concept relative to the concept: THE PARTICULAR ANIMAL USED BY LAWRENCE TO CROSS THE ARABIAN DESERT. Moreover, it is common in cognitive science to refer to language, mindreading, and social exchange as domains. Are these things encompassing in the same sense and to the same degree as ANIMAL? To these difficulties, some researchers define domain as a sets of principles. This won’t help. We have principles underlying our knowledge of camels, as well as principles underlying our knowledge of animals. I see no escape. If we drop the weak definition of “domain” (domain = subject matter), we still find ourselves with definitions that are vague or insufficiently restrictive.

Things are slightly better with “specificity.” On a strong reading, “specific” means “exclusively dedicated.” To say that modules are domain specific is to say that they are exclusively dedicated to their subject matter. This is a useful explanatory construct, and it may be applicable to certain mental systems. Consider the columns of cells in primary visual cortex that are used to detect edges. These cells may be dedicated to that function and nothing else. Perhaps modules are supposed to be like that.

There is still some risk of triviality here. We can show that any collection of rules and representations in the mind-brain is dedicated by simply listing an exhaustive disjunction of everything that those rules and representations do. To escape triviality, we want to rule out disjunctive lists of functions. We say that systems are domain specific when the domain can be specified in intuitively coherent way. Let’s assume for the sake of argument that this requirement can be made more precise. The problem is that alleged examples of modules probably aren’t domain specific in this strong sense.

Consider vision. Edge detectors may be domain specific, but other resources used for processing visual information may be more general. For example, the visual system can be recruited in problem solving, as when one uses imagery to estimate where a carton of milk can squeeze into a crammed refrigerator. Some of our conceptual knowledge may be stored in the form of visual records. We know that damage to visual areas can disrupt conceptual competence (Martin & Chao, 2001). I have also noted that, when people lose their sense of sight, areas once used for vision get used for touch. Visually perceived stimuli also generate activity in cells that are bimodal. The very same cells are used by the touch system and the auditory system. If we excluded rules and representations that can be used for something other than deriving information from light, the boundaries of the “visual system” would shrink considerably. At the neural level of description, it is possible that only isolated islands of cells would remain. This would be a strange way to carve up the mind. One of the important things about our senses is that

they can moonlight. They can help each other out and they can play a central role in the performance of cognitive tasks. Vision, taken as a coherent whole, is not domain specific in the strong sense, even if it contains some rules and representations that are.

Similar conclusions can be drawn for language. I have said that language may share resources with systems that serve other functions: pattern recognition, muscle control, and so on. Broca's area seems to contain mirror neurons, which play a role in the recognition of manual actions, such as pinching and grasping (Heiser et al. 2003). Wernicke's area seems to contain cells that are used in the categorization of non-linguistic sounds (Saygin et al. 2003). Of course, there may be some language-specific rules and representation *within* the systems that contribute to language. Perhaps the neurons dedicated to conjugating the verb "to be" have no nonlinguistic function. Such highly localized instances of domain specificity will offer little comfort to the defender of modularity. They are too specific to correspond to modules that have been proposed. Should we conclude that there is a module dedicated to the conjugation of each irregular verb?

There is relatively little evidence for large-scale, if we use domain specificity as the criterion. It is hard to find systems that are exclusively dedicated to broad domains. Vision and language systems are not dedicated in the strong sense, and the same is true for other alleged modules. Consider mindreading, which clearly exploits domain general capacities. I noted above that mindreading is correlated with language skills. Hale and Tager-Flusberg (2003) found that preschoolers who failed the false belief task were more likely to succeed after receiving training in sentential complement clauses. They went from 20% correct in attribute false beliefs to over 75% correct. Mindreading also depends on working memory. Performance in attributing false beliefs is impaired in healthy subjects when they are given an unrelated working memory task (McKinnon & Moscovitch, unpublished). In neuroimaging studies, mindreading is shown to recruit language centers in left frontal cortex, visuospatial areas in right temporal-parietal regions, the amygdala, which mediates emotional responses, and the precuneus, which is involved in mental image inspection and task switching. In short, mindreading seems to exploit a large network of structures all of which contribute to many other capacities.

This seems to be the general pattern for alleged modules. The brain structures involved in mathematical cognition are also involved in language, spatial processing, and attention (Dehaene, 1997; Simon, 1997). Folk physics seems to rely on multi-object attention mechanisms (Scholl & Leslie, 1999). Moral judgment recruits ordinary emotion centers (Greene & Haidt, 2002).

For all I have said, alleged modules may have domain specific *components*. Perhaps these systems use some proprietary rules and representations. But they don't seem to be proprietary throughout. Therefore, domain specificity cannot be used to trace the boundaries around the kinds of systems that modularists have traditionally discussed.

6. Inaccessibility and Encapsulation

The final two properties on Fodor's list are closely linked. Modules are said to be inaccessible and encapsulated. That means, they don't let much information out and they don't let much information in. Fodor thinks the latter property is especially important.

Carruthers places emphasis on both encapsulation and inaccessibility. I think neither property is especially useful in carving up the mind.

Let's begin with inaccessibility. Fodor claims that systems outside a module have no access to the internal operations within that module. This seems plausible introspectively. I have no introspective access to how my visual system achieved color constancy or how my syntax system parses sentences. Nisbett and Wilson (1977) have shown that human judgment is often driven by processes that operate below the level of consciousness. Does this confirm that operations within modules are inaccessible? No; it shows only that we lack conscious access. It tells us nothing about whether operations within unconscious mental systems are accessible to other unconscious systems. For all we know, there may be extensive accessibility below the level of awareness.

This is where Carruthers comes in. He has a principled argument for the conclusion that mental systems are by and large inaccessible. He says that, in order for one system to access information in another, the first system would need to represent how the other system works. But that means it would need to represent all the rules and representations of that other system. This would defeat the purpose of dividing the mind into separate systems, and it would lead to a combinatorial explosion. Therefore, most systems must be inaccessible to each other.

I am not persuaded by this argument. It rules out the view that all systems are *completely* accessible to all others, but it does nothing to refute the possibility that *some* systems have *some* access to others. For example, conceptual systems might have access to syntactic trees, but lack access to subtle transformation rules used in deriving those trees. Limited accessibility would not lead to a combinatorial explosion, and it might be useful for some systems to have an idea what other systems are doing. By analogy, the President cannot attend every cabinet meeting, but it would help him to have some idea of how cabinet members reached any given decision.

Let me turn from inaccessibility to encapsulation—the final item on Fodor's list and, for him, the most important. Fodor tries to prove that perceptual systems are modular by appealing to perceptual illusions. This interesting thing about illusions is that they persist even when we know that we are being deceived. The two lines in the Müller-Lyre illusion appear different in length even though we know they are the same. If perception were not encapsulated, then the illusion would go away as soon as the corrective judgment is formed. Belief would correct experience.

Fodor's argument is flawed. There are competing explanations for why illusions persist. One possibility is that perception always trumps belief when the two come into conflict. Such a trumping mechanism would be advantageous, because, otherwise, we could not use experience to correct our beliefs. The trumping mechanism is consistent with the hypothesis perception is not encapsulated. Beliefs may be able to affect perception when the two are not in conflict. To test between trumping and encapsulation, we need to consider such cases. Consider ambiguous figures. Verbal cueing can lead people to alter their experience of the duck-rabbit. Likewise, we can electively experience a Necker cube as facing right, facing left, or as a gemstone facing directly forward. In paintings that convey depth by scale, we can see figures in the distance as far away or we can see them as tiny people floating in the foreground.

Other examples of top-down effects are easy to generate. For example, expectations can lead us to experience things that aren't there. If you are waiting for a

visitor, every little sound may be mistaken for a knock on the door. Or consider visual search: when looking for a Kodak film carton, small yellow objects pop out in that visual field. The most obvious case of top-down influence is mental imagery. Cognitive states can be used to actively construct perceptual representations (Kosslyn et al. 1995). This makes sense of the neuroanatomy: there are dense neural pathways from centers of higher brain function into perception centers.

There is also evidence for top-down-effects in language processing. Defenders of modularity would have us believe that language divides into a number of modular subsystems, including syntax, semantics, and phonology. These subsystems are alleged to be impervious to each other, but there is empirical evidence to the contrary. For example, in the phenomenon of phoneme restoration, subjects are presented with sentences containing deleted phonemes, but, rather than hearing an acoustic gap, the missing phoneme is filled-in. Importantly, the phoneme that is heard is determined by the semantic interpretation of the sentence (Warren & Warren, 1970). If subjects hear, "The _eel is on the axle," they experience a "w" sound in the gap. If they hear, "The _eel is on the orange," they experience a "p" sound.

There is also evidence that syntax can be affected by conceptual knowledge. Marslen-Wilson and Tyler (1987) showed that conceptual factors exert highly specific influences on sentence completion, and they do so at the same speed as lexical factors. In one experiment, subjects are given the following story: "As Philip was walking back from the shop he saw an old woman trip and fall flat on her face in the street. She seemed unable to get up again." The story then continues with one of two sentence fragments: either "He ran toward..." or "Running towards..." In both fragments, the appropriate next word is "her," but in the first case that choice is determined *lexically* by the prior pronoun in the sentence ("he") and in the second case that choice is determined *conceptually* (we know that people cannot run when they are lying down). Remarkably, subjects are primed to use the word "her" equally fast in both conditions. If lexical processing were encapsulated from conceptual processing, one would expect lexically determined word choices to arise faster. These results imply that formal aspects of language are under immediate and constant influence of general world knowledge.

Thus far, I have been talking about top-down influences on input systems. There is also evidence that input systems can speak to each other. This is incompatible with encapsulation, because a truly encapsulated system would be insulated from *any* external influence. Consider some examples. First, when subjects hear speech sounds that are inconsistent with observed mouth movements, the visual experience systematically distorts the auditory experience of the speech sounds (McGurk & McDonald, 1976). Second, Ramachandran has developed a therapeutic technique for treating phantom limb pain, in which amputees use a mirror reflection to visually relocate an intact limb in the location of a missing limb; if they scratch or soothe the intact limb, the discomfort in the phantom subsides (Ramachandran et al. 1995). Third, sound can give rise to touch illusions: hearing multiple tones can make people feel multiple taps, when there has been only one (Hötting & Röder, 2004). Finally, people with synesthesia experience sensations in one modality when they are stimulated in another; for example, some people see colors when they hear sounds, and others experience shapes when they taste certain flavors. All these examples show that there can be direct and content-specific cross-talk between the senses.

The empirical evidence suggests that mental systems are not encapsulated. But the story cannot end here. There is also a principled argument for encapsulation, which is nicely presented by Carruthers. It goes like this: mental processes must be computationally tractable, because the mind is a computer, and mental processes are carried out successfully in a finite amount of time; if mental processes had access to all the information stored in the mind (i.e., if they were not encapsulated), they would not be tractable (merely checking consistency against a couple hundred beliefs would take billions of years); therefore, mental processes are encapsulated.

Carruthers recognizes that there is a major flaw in this argument. According to the second premise, mental processes would be intractable if they had access to all the information stored in the mind. This is actually false. Computational systems can sort through stupendously large databases at breakneck speed. The trick is to use “frugal” search rules. Frugal rules are ones that radically reduce processing load by exploiting simple procedures for selecting relevant items in the database. Once the most relevant items are selected, more thorough processing of those items can begin. Psychologists call such simple rules “heuristics” (Kahneman et al. 1982). There is overwhelming evidence that we make regular use of heuristics in performing cognitive tasks. For example, suppose you want to guess which of two cities is larger, Hamburg or Mainz. You could try to collect some population statistics (which would take a long time), or you could just pick the city name that is most familiar. This Take the Best strategy is extremely easy and very effective; it is even a good way to choose stocks that will perform well in the market (Gigerenzer et al., 1999). With heuristics, we can avoid exhaustive database searches even when a complete database is at our disposal. There are also ways to search through a colossal database without much cost. Internet search engines provide an existence proof (Clark, 2002). Consider Google. A Google search on the word “heuristic” sorts through over a billion web pages in 0.18 seconds, and the most useful results appear in the first few hits. Search engines look for keywords and for webpages that have been frequently linked or accessed. If we perform the mental equivalent of a Google search on our mental files, we should be able to call up relevant information relatively quickly. The upshot is that encapsulation is not needed for computationally tractability.

At this point, one might expect Carruthers to the assumption that mental systems are encapsulated. Instead, he draws a distinction between two kinds of encapsulation. Narrow-scope encapsulation occurs when most of the information held in the mind is such that a system can’t be affected by that information in the course of processing. This is the kind of encapsulation that Fodor attributes to modules, and it is what Carruthers rejects when he appeals to heuristics. It is possible that any item of information is such that a system could be affected by it. But Carruthers endorses wide-scope encapsulation: systems are such that they can’t be affected by most of the information held in the mind at the time of processing. This seems reasonable enough. If every item in the mind sent inputs to a given system simultaneously, that system would be overwhelmed. So, I accept “wide-scope encapsulation.” But “wide-scope encapsulation” is not really encapsulation at all. “Encapsulation” implies that one system cannot be accessed by another. “Wide-scope encapsulation” says that all systems are accessible; they just aren’t accessed all at once. Carruthers terminological move cannot be used to save the

hypothesis that mental systems are encapsulated. In recognizing the power of heuristic search, he tacitly concedes that the primary argument for encapsulation is unsuccessful.

I do not want to claim that there is *no* encapsulation in the mind. It is possible that some subsystems are impervious to external inputs. I want to claim only that there is a lot of cross-talk between mental systems. If we try to do mental cartography by drawing lines around the few subsystems that are encapsulated, we will end up with borders that are not especially helpful. Encapsulation it is not sufficiently widespread to be an interesting organizing principle.

9. Conclusion: Decomposing Modularity

Throughout this discussion, I have argued that Fodor's criteria for modularity do not carve out interesting divisions in the mind. Systems that have been alleged to be modular cannot be characterized by the properties on Fodor's list. At best, these systems have *components* that satisfy some of Fodor's criteria. There is little reason to think that these criteria hang together, and, when considered individually, they apply to a scattered and sundry assortment of subsystems. It is grossly misleading to say that the mind is modular. At best, the mind has a smattering of modular parts.

That does not mean that the mind is a disorganized mash. At the outset, I said that modularity is not equivalent to functional decomposition. The mind can be described as a network of interconnected systems and subsystems. We can represent the mental division of labor using flowcharts whose units corresponding to functionally distinguished components that carry out subroutines and contribute, in their limited way, to the greater whole. My goal has been to criticize a specific account of what the functional units in the mind are like. The functional units need not be fast, automatic, innate, shallow, or encapsulated. Some of the components may be dedicated to a single mental capacity, but others may serve a variety of different capacities. It is possible that no component in the mind exhibits the preponderance of properties on Fodor's list.

Some defenders of modularity are committed to nothing more than functional decomposition. They reject Fodor's list and adopt the simple view that the mind is a machine with component parts. *That* view is uncontroversial. Massive modularity sounds like a radical thesis, but, when the notion of modularity is denatured, it turns into a platitude. Of course central cognition has a variety of different rules and representations. Of course we bring different knowledge and skills to bear when we reason about the social world as opposed to the world of concrete objects. Of course it is possible for someone to lose a specific cognitive capacity without losing every other cognitive capacity. Controversy arises only when functional components are presumed to have properties on Fodor's list.

I think the term "modularity" should be dropped because it implies that many mental systems are modular in Fodor's sense, and that thesis lacks support. Cognitive scientists should continue to engage in functional decomposition, but we should resist the temptation to postulate and proliferate modules.

References

- Baillargeon, R., L. Kotovsky, & A. Needham (1995). The acquisition of physical knowledge in infancy. In D. Sperber, D. Premack, and A. J. Premack, eds., *Causal cognition: A multidisciplinary debate*. New York: Oxford University Press.
- Baron-Cohen, S. (1996). *Mindblindness: An essay on autism and theory of mind*. Cambridge, MA: MIT Press.
- Bates, E. (1994). Modularity, domain specificity and the development of language. In D.C. Gajdusek, G.M. McKhann, & C.L. Bolis, (Eds.), *Evolution and neurology of language. Discussions in Neuroscience*, 10, 136-149.
- Bates, E., Dale, P., & Thal, D. (1995). Individual differences and their implications for theories of language development (pp. 96-151). In P. Fletcher and B. MacWhinney (Eds.), *Handbook of child language*. Oxford: Blackwell.
- Bates, E., Marangolo, P., Pizzamiglio, L., Devescovi, A., Ciurli, P., & Dick, F. (2000). Linguistic and nonlinguistic priming in aphasia. *Brain and Language*, 76, 62-69.
- Bishop, D. V. (1992). The underlying nature of specific language impairment. *Journal of Child Psychology and Psychiatry*, 33, 3-66.
- Caramazza, A. & Mahon, B. Z. (2003). The organization of conceptual knowledge: The evidence from category-specific semantic deficits. *Trends in Cognitive Science*, 7, 354-361.
- Chen, R., Cohen, L. G., Hallett, M. (2002). Nervous system reorganization following injury. *Neuroscience*, 111, 761-773.
- Clark, A. (2002). Local associations and global reason: Fodor's frame problem and second-order search. *Cognitive Science Quarterly*, 2, 115-140.
- Cosmides, L. & J. Tooby (1994). Origins of domain specificity: The evolution of functional organization. In L. A. Hirschfeld & S. A. Gelman (eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 85-116). New York: Cambridge University Press, 85-116.
- Dehaene, S. (1997). *The number sense*. New York: Oxford University Press.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fodor, J. (2000). *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press.
- Garfield, J. L., Peterson, C. C. & Perry, T. (2001). Social cognition, language acquisition and the development of the theory of mind. *Mind & Language*, 16, 494-541.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Greene, J. & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Science*, 6, 517-523.
- Hale C. M. & Tager-Flusberg H. (2003). The influence of language on theory of mind: A training study. *Developmental Science*, 6, 346-359.
- Heiser, M., Iacoboni, M., Maeda, F., Marcus, J. & Mazziotta, J.C. (2003). The essential role of Broca's area in imitation. *European Journal of Neuroscience*, 17, 1123-1128.
- Hood, B., Carey, S., & Prasada, S. (2000). Predicting the outcomes of physical events: Two-year-olds fail to reveal knowledge of solidity and support. *Child Development*,

71, 1540-1554.

- Hötting, K. & Röder, B. (2004). Hearing Cheats Touch, but Less in Congenitally Blind Than in Sighted Individuals. *Psychological Science*, 15, 60-64.
- Insel, T. R., O'Brien, D. J., & Leckman, J. F. (1999). Oxytocin, vasopressin, and autism: is there a connection? *Biological Psychiatry*, 45, 145-157.
- Jacobs, R. A. (1999). Computational studies of the development of functionally specialized neural modules. *Trends in Cognitive Science*, 3, 31-38.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Kosslyn, S. M., Thompson, W. L., Kim, I. J., & Alpert, N. M. (1995). Topographical representations of mental images in primary visual cortex. *Nature*, 378, 496-498.
- Lashley, K. (1950). In search of the engram. *Symposia of the Society for Experimental Biology*, 4, 454-482.
- Leslie, A. M. (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*, (pp. 119-148). New York: Cambridge University Press.
- Lillard, A. (1998). Ethnopsychologies: Cultural variations in theories of mind. *Psychological Bulletin*, 123, 3-32.
- Marslen-Wilson, W. & Tyler, L. (1987). Against modularity. In J. L. Garfield (Ed.), *Modularity in knowledge representation and natural-language understanding* (pp. 37-62). Cambridge, MA: MIT Press.
- Martin, A., and Chao, L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion In Neurobiology*, 11, 194-201.
- McGurk, H. & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264, 746-748.
- McKinnon, M. & Moscovitch, M. (unpublished). Domain-general contributions to social reasoning: Perspectives from aging and the dual-task method. Manuscript, University of Toronto.
- Needham, A. and Baillargeon, R. (1993) Intuitions about support in 4.5-month-old infants. *Cognition* 47, 121-148.
- Nisbett, R., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231-259.
- Pinker, S. (1997). *How the Mind Works*. New York : Norton.
- Plaut, D. C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, 17, 291-321
- Poepfel, D. (1996). A critical review of PET studies of phonological processing. *Brain and Language*, 55, 317-351.
- Prinz, J. J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22, 253-336.
- Quartz, S. R., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioural and Brain Sciences*, 20, 537-596.
- Ramachandran, V. S., Rogers-Ramachandran, D., Cobb, S. (1995). Touching the

- phantom limb. *Nature*, 377, 489-90.
- Samuels, R. (1998). Evolutionary psychology and the massive modularity hypothesis. *British Journal for the Philosophy of Science*, 49, 575-602.
- Saygin, A.P., Dick, F., Wilson, S.M., Dronkers, N.F. & Bates, E. (2003) Neural resources for processing language and environmental sounds: Evidence from aphasia. *Brain*, 126(4), 928-945.
- Scholl, B.J., & Leslie, A.M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In (Eds.), E. Lepore & Z. Pylyshyn, *What is Cognitive Science?* (pp. 26–73). Oxford: Blackwell.
- Seung, H-K. & Chapman, R. S. (2000). Digit span in individuals with Down syndrome and typically developing children: Temporal aspects. *Journal of Speech, Language, and Hearing Research*, 43, 609-620.
- Simon, T. J. (1997). Reconceptualizing the origins of number knowledge: A non-numerical account. *Cognitive Development*, 12, 349–372.
- Slater, A., Morison, V., Somers, M., Mattock, A., Brown, E., & Taylor, D. (1990). Newborn and older infants' perception of partly occluded objects. *Infant Behavior and Development*, 13, 33–49.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (eds), *Mapping the mind: Domain specificity in cognition and culture* (pp. 29-67). New York: Cambridge University Press.
- Stone, V. E., Cosmides, L., Tooby, J., Kroll, N. & Knight, R.T. (2002). Selective impairment of reasoning about social exchange in a patient with bilateral limbic system damage. *Proceedings of the National Academy of Sciences*, 99, 11531-11536.
- Tallal, P., Miller, S. L., Bedi, G., Byma, G., Wang, X., Nagarajan, S. S., Schreiner, C., Jenkins, W. M., & Merzenich, M. M. (1996). Language comprehension in language-learning impaired children improved with acoustically modified speech. *Science*, 271, 81–84.
- Thomas, M. S. C. & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioural and Brain Sciences*, Vol.25, No.6, 727-788.
- Uttal, W. R. (2001). *The new phrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press.
- Van Giffen, K. & Haith, M. M. (1984). Infant response to Gestalt geometric forms. *Infant Behavioral Development*, 7, 335-46 .
- Vargha-Khadem, F., Watkins, K., Alcock, K., Fletcher, P., & Passingham, R. (1995). Praxic and nonverbal cognitive deficits in a large family with a genetically transmitted speech and language disorder. *Proceedings of the National Academy of Science*, 92, 930–933.
- Vinden, P. G. (1996). Junin Quechua children's understanding of mind. *Child Development*, 67, 1701-1716.
- Vinden, P. G. (1999). Children's understanding of mind and emotion: A multi-culture study. *Cognition and Emotion*, 13, 19-48
- Warren, R. M., & Warren, R. P. (1970). Auditory illusions and confusions. *Scientific American*, 223, 30-36.