What Is Domain Specificity (and Why Does It Matter)?

Muhammad Ali Khalidi (khalidi@yorku.ca)
Department of Philosophy and Cognitive Science Program,
York University, 4700 Keele Street
Toronto, ON M3J 1P3 Canada

Abstract

The distinction between domain specificity and domain generality is widespread in cognitive science. Yet, the difference between the two types of cognitive capacities has rarely been made in a principled manner. Moreover, some of the examples that are put forward to illustrate it in the literature are either spurious or misleading. In this paper, I use a number of examples to determine what domain specificity is, and just as importantly, what it is not. A domain-specific cognitive system is one that is in principle generalizable, but which the cognizer does not extend to cases that the system did not originally evolve to deal with.

Keywords: domain specificity; modularity; innateness; animal cognition.

Introduction

There are many contexts in cognitive science in which it is useful to distinguish domain-specific cognitive capacities from domain-general ones. For instance, according to many evolutionary psychologists, human cognition consists largely of domain-specific cognitive capacities, and this feature of our cognitive makeup provides evidentiary support for the pervasive influence of evolutionary processes on the formation of the human mind (Carey & Spelke 1994, Cosmides & Tooby 1994). By contrast, according to other researchers, domain-general cognitive abilities are the norm in the human mind and are what distinguish human cognition from that of most other animals (Samuels 1998, Fodor 2000). Yet, the distinction between the two kinds of capacities is not easily drawn and it is often drawn erroneously by the researchers who aim to make it. After identifying some misleading examples of domain specificity, I will put forward some better examples of the phenomenon drawn from the literature. Then, I will use these examples to try and spell out a more satisfactory account of the phenomenon of domain specificity. Finally, I will draw on this new understanding to try to shed light on the significance of this concept and its importance for cognitive science.

Domain Specificity and Its Confounds

Domain specificity is a feature of cognitive capacities that is often associated with several other such features, notably: modularity, innateness, and brain localization. In this section, I will examine the connections that may or may not exist between domain specificity and these other features, in order to gain a preliminary understanding of the concept of domain specificity.

By virtue of the way in which modularity was initially defined by Fodor (1983), there is a strong link between modularity and domain specificity. Indeed, it follows from Fodor's account that domain specificity is one of the defining features of modularity, and therefore that all modular cognitive capacities are domain-specific.[1] Of course, a case might be made for rejecting this definition on the grounds that it is unwarranted by the empirical facts or otherwise detrimental to cognitivist research, but I will not try to make the case, nor do I think that the case ought to be made. I will simply follow Fodor and much subsequent theorizing in accepting it. Hence, I take it as uncontroversial that domain specificity is a necessary feature of modularity, though the two features ought not to be conflated, since the former is subsumed by the latter.

The case is more complicated when it comes to innateness. Although there is also a widespread assumption that there is a link between innateness and domain-specificity, there is no convincing reason for inferring such a link. It is neither the case that all innate cognitive capacities are domain-specific, nor that all domain-specific cognitive capacities are innate. To illustrate, human beings may have an innate cognitive capacity for associative learning that may be entirely domain-general. Conversely, there may be certain domain-specific cognitive abilities that are not innate but mainly learned, such as chess-playing ability. However, despite the lack of a direct link between the two features of cognition (innateness and domain specificity), a case could be made for an indirect link between them. It has been argued that when it comes to domain-specific abilities, it is easier to tell whether and to what extent they are innate or not (Khalidi 2001). In other words, the link between the two features is epistemic or evidential rather than intrinsic, since we can more easily gauge the amount of explicit learning or relevant experience in the case of domain-specific cognitive capacities than in the case of domain-general ones. This is so because it is easier to rule out relevant sources of information in the former case than in the latter.

As for the link between domain specificity and brain localization, this is also widely made, as is the link between

---

[1] In addition to being domain-specific, according to Fodor (1983) modular cognitive capacities are supposed to: 2) process items automatically and in a mandatory manner, 3) be inaccessible to consciousness, 4) be fast, 5) be cognitively impenetrable (e.g. resistant to being unlearned), 6) process "shallow" or highly salient features, 7) have fixed neural architecture, 8) have specific breakdown patterns (as in aphasia, agnosia), and 9) have fixed ontogeny (standard pace and sequence of development).

brain localization and modularity (which, as seen above, subsumes domain specificity). However, there does not seem to be a cogent reason for making either link. For instance, there are good grounds for thinking that various cognitive capacities are modular (and hence domain-specific) even though they are not localized in one region of the brain, indeed even though they are scattered across a range of brain regions. Modularity and domain-specificity pertain largely to the functioning of a cognitive capacity rather than its neural manifestation, so there is limited scope for inferring brain localization from either of these phenonema.

In distinguishing domain specificity from other features of human cognition, I have relied on an implicit preliminary understanding of the phenomenon. As the term implies, what it is for a cognitive capacity to be domain-specific is for it to pertain to a single domain or to a restricted range of domains, and more importantly, for it not to be generalizable to other domains. Moreover, this last proviso highlights the importance of restricting domain specificity to aspects of cognition that are in principle *generalizable* across domains, although they are not in fact *generalized*. In other words, it would be vacuous to describe as domain-specific some cognitive system that is not even in principle generalizable. For example, a body of information pertaining to some domain or another is not generalizable, since its subject matter is in principle restricted to a certain domain. By contrast, a rule that is deployed by a cognizer in one domain but that *could* be deployed in another domain is in principle generalizable. Hence, rather than domain-specific capacities it may be more fruitful to talk about domain-specific *rules* or *principles*. These are the kinds of cognitive structures that can be generalized across domains and that are therefore candidates for being domain specific, though they may not be the only cognitive structures that can be so generalized. This point will be developed further in subsequent sections.

One issue that might be raised here concerns the nature and scope of a domain. If we do not explicitly specify what a domain consists in, we may open ourselves to the following objection. Suppose it is claimed that some linguistic rule *R*, which is part of a subject's cognitive repertoire, is domain-specific since it only pertains to the domain of language. Now suppose someone objects to this on the grounds that language itself is not a single domain but rather that it comprises various domains, say syntax, semantics, pragmatics, and so on. Hence, this objector might contend that the rule is in fact a domain-general one since it ranges over several domains. This example suggests that it may always be open for someone to claim that a putative domain-specific rule is in fact domain-general unless we have some principled way of delimiting the scope of a domain. How could we define domain specificity without specifying the scope of a domain? In response, I would argue that it is not enough to show that a rule does in fact apply to what are supposed to be different domains. Rather, to demonstrate true domain generality with respect to some

rule *R* used by some cognizer *C*, it should be shown that *C* has the ability to deploy *R* in some previously unencountered situation or to apply it to some new cases. In practice, it does not seem to pose much of a threat to our account of domain specificity if we do not have a principled way of distinguishing domains. Typically, it will suffice to show that a cognitive mechanism is truly domain-general if we can show that it can be deployed in unfamiliar contexts that are new to the cognizer. Conversely, it suffices to show that a cognitive mechanism is domain-specific if we can demonstrate that when one attempts to apply it to new cases, the mechanism either does not function or gives systematically erroneous answers. This point will be justified further in due course in discussing what constitutes a new case or a new domain.

Preliminary Examples

Spurious examples of domain specificity, or genuine examples of domain specificity that are misleadingly described, are not difficult to find in cognitive science. Two such instances stand out in Cosmides and Tooby (1994), in the context of an argument that domain-specific cognitive mechanisms are more efficient and adaptive that domain-general ones. Cosmides and Tooby (1994, 90) state: "A woman who used the same taste preference mechanisms in choosing a mate that she used to choose nutritious foods would choose a very strange mate indeed, and such a design would rapidly select itself out." Clearly, something has gone wrong here: a rule that chooses mates on the same basis as foods is not a domain-general rule but one that is based on an error. If we have a cognitive mechanism that enables us to recognize edible foods, it should only be sensitive to perceptual stimuli that are plausible candidates for food. Otherwise, it is not a genuine food-preference cognitive mechanism. Hence, this is not a plausible example of a domain-general mechanism being inferior to a domain-specific one. In describing the example, Cosmides and Tooby seem to be trading on the ambiguity of the term "taste", which can either refer to gustatory discrimination or a broader notion of discrimination, which could comprise mate selection.

A case of domain specificity that is misleadingly described can be found in another example given by Cosmides and Tooby (1994). Although the following example is a genuine case of domain specificity, it is improperly contrasted with a putative case of domain generality. By examining it, we can get a better handle on the phenomenon of domain specificity. Cosmides and Tooby explicate the well-known example of the alarm calls of vervet monkeys, who give three different calls in response to three different kinds of predators (leopard, eagle, and snake), leading conspecifics to take three different types of evasive action (respectively, climbing a tree, looking up or diving into bushes, and standing on hind legs and looking into the grass). In this case, they state: "A single, general-purpose alarm call (and response system) would be less effective because the recipients of the call

would not know which of the three different and incompatible evasive actions to take" (Cosmides & Tooby 1994, 89-90). The problem here is not there could not be a general-purpose alarm system; there clearly could. But a general-purpose alarm system is not one that would issue the same call for every predator. That would be a system that fails to discriminate among different stimuli. Rather, an all-purpose alarm system would be one akin to the human linguistic alarm system, which issues a different linguistic warning in the case of different predators. There are clearly certain advantages to such a system, since it is capable of handling a much wider range of predators ("Lion!", "Hawk!", "Stampede of buffalo!") and of being made more precise in various ways ("Tiger to the right", "Hyena to the northwest", "Human with weapon right behind you", etc.). However, it may also involve certain disadvantages, since given the diversity of inputs and outputs, it may take more processing time to issue the correct alarm, there may be more opportunity for error in both transmission and reception, and the evasive action involved may have to be figured out from scratch by the respondent once the alarm is sounded. Determining which of these two alarm systems, the domain-specific vervet system or the domain-general human system, is more efficient and adaptive is not an easy matter. It will clearly depend on various contingencies such as the nature of the environment, and Cosmides and Tooby may ultimately be right that in certain circumstances a domain-specific system may be superior to a domain-general one. But they have not made the appropriate contrast between domain specificity and domain generality.

This example is instructive since, once modified in the way that I have just done, it provides a fairly clear contrast between a domain-specific and a domain-general cognitive mechanism. In the following section I will attempt to give a more satisfactory formulation of domain specificity and illustrate it with better examples drawn from the literature.

Domain Specificity Revised

To refine our understanding of domain specificity, it is useful to build on the example mentioned in the previous section and attend to its instructive features. The first feature that can be gleaned from the vervet monkey alarm call system is that a cognitive mechanism for alarm calls is at least in principle generalizable. That is to say, even though the vervet alarms are only issued for a small set of specific predators, it is easy to conceive of an alarm system that would extend to other predators. Hence, it seems safe to conclude that for one to speak meaningfully of a domain-specific cognitive system, that system must have the following feature:

1) A *domain-specific* cognitive system is one that is in principle generalizable to new domains.

This condition may appear vacuous, but it is designed to rule out cognitive systems that consist of a "database" rather than rules or principles, as mentioned above. Domain specificity, to be meaningful, must be a feature of the cognitive capacity rather than a feature of the subject matter. Though this point may seem obvious, the attribute of domain specificity is often conferred on bodies of knowledge possessed by subjects that are not obviously generalizable, such as knowledge of animate as opposed to inanimate domains (e.g. Caramazza & Shelton 1998).

The first proposed feature of domain specificity makes reference to "new domains," which is a notion that needs further explication and justification. The new domain involved need not be what we might regard as an entirely disparate area of inquiry. In the case of the vervets, the original domain is something like: predators commonly encountered by vervets in the wild. It is *in principle* generalizable to include the new domain: all predators, or even, all threats. The vervet alarm system is domain-specific because it fails to generalize to these new stimuli. But these new stimuli do not, strictly speaking, have to be drawn from what we would normally consider to be another domain, such as a new sensory modality or a new area of inquiry. At this point, it might be asked, by virtue of what are they to be considered genuinely new stimuli? They must at least be stimuli that the cognizer has not encountered before. But that condition is surely too weak, since the domain-specific vervet alarm system clearly generalizes to new exemplars of leopards, eagles, and snakes, which the individual has not encountered before, indeed ones which perhaps no vervet monkey has encountered before. Rather, in this context, new stimuli are ones that the system was not originally designed to cope with. This is admittedly a vague formulation and brings in thorny evolutionary considerations concerning the *proper function* of an evolved trait (Millikan 1989, Neander 1991). Though it is not always easy to determine what the proper function of a cognitive system is, some reference to it seems inevitable, since cognitive systems have evolved to fulfill a certain purpose and their generalizability consists in part in being able to extend beyond that purpose to cases that they were not designed to cope with, or ones that are not normally encountered in the environment in which the system evolved.[2] Hence, I propose that the second crucial feature of a domain-specific cognitive system is as follows:

2) A *domain-specific* cognitive system is one that systematically fails to yield a correct result in the case of stimuli that the system did not evolve to deal with.

The aptness of this second feature can be further justified by reflecting on appropriate examples from the literature. One such case is provided by Cheney and Seyfarth (1985, 197), who describe the domain-specificity of certain cognitive capacities in vervet monkeys, as follows: "Within the social group, the behavior of monkeys suggests an understanding of causality, transitive inference, and the notion of reciprocity. Despite frequent opportunity and often

---

[2] A similar conclusion has been reached by Boyer and Barrett (2005, 98), who write: "The domain of operation of the system is best circumscribed by evolutionary considerations."

strong selective pressure, however, comparable behavior does not readily emerge in dealings with other animal species or with inanimate objects." In this example, both features outlined above are clearly in evidence. First, the principle of causality and the rule of transitivity clearly have application outside the realm of social interaction with conspecifics. The transitivity rule can be used to infer hierarchy relations among monkeys (if *A* ranks higher than *B* and *B* ranks higher than *C*, then *A* ranks higher than *C*) but it can also be used to infer information about size, quantity, and other matters (if object *A* is larger than object *B* and *B* is larger than *C*, then *A* is larger than *C*). However, despite the clear applicability of this rule to domains that go beyond social interactions with conspecifics, Cheney and Seyfarth claim that vervets do not so apply the rule. Second, it is clear that vervets do not use the rule of transitivity on other species or inanimate objects simply because they evolved the rule to deal with the restricted domain of social interaction with conspecifics, which may have been a more pressing adaptive problem. In this case, it may seem obvious that interactions with other animal species and with inanimate objects constitute genuinely new domains. There may not appear to be a need to use the second feature of domain specificity to justify the judgment that it does not generalize to genuinely new domains. But even though it may not play this role in this case, there are other cases in which the second condition is necessary to enable us to make a judgment concerning domain specificity.

Some recent studies of cognition in non-human animals have debated whether animals have the cognitive capacity to teach others. Thornton and McAuliffe (2006) demonstrate that mature meerkats provide young conspecifics with specimens of their usual prey, scorpions, that are either dead, disabled (stings removed), or intact. Which of these three types of scorpion is provided depends on the perceived age of the young meerkat. Younger pups are provided with disabled but alive scorpions and this provides them with the opportunity to learn how to kill the scorpion without being exposed to the possibility of a sting. Thornton and MacAuliffe (2006, 228) conclude that "the provisioning behavior of meerkat helpers constitutes a form of 'opportunity teaching' in which teachers provide pupils with opportunities to practice skills, thus facilitating learning." Moreover, they support their findings by relying on a definition of teaching derived from Caro and Hauser (1992), according to which an individual is a teacher if it modifies its behavior in the presence of a naïve observer, at some cost to itself, and as a result allows the observer to acquire knowledge or skills. In this case, it appears more difficult to rule decisively that the teaching is domain-specific rather than domain-general, since as Csibra (2007, 96) argues: "the opportunity teaching that has been demonstrated in meerkats does result in the acquisition of a generalizable skill: it not only provides youngsters with food but also 'teaches' them how to kill scorpions." However, despite the apparent (limited) generalizability of this skill, what rules it out as a genuinely domain-general skill is that it appears to be restricted to behaviors designed to facilitate preying on scorpions. Moreover, this is likely to be the function for which this behavior was evolved. Unless one can demonstrate otherwise, it is safe to conclude that this is a domain-specific capacity. (Whether that rules it out as a genuine case of teaching is another matter.) Hence, the second condition on domain specificity enables us provisionally to decide that this is indeed a domain-specific cognitive capacity.

Further Evidence

So far, the examples I have considered derive primarily from studies in evolutionary psychology and comparative cognition. But the concept of domain specificity has also had considerable influence in cognitive neuroscience and developmental psychology. I will now consider whether the notion as I have characterized it can be pressed into service in other areas of cognitive science.

There is a well-established body of evidence indicating the existence of category-specific semantic deficits in a range of patients with brain lesions and other neural abnormalities. However, the correct interpretation of this evidence remains a source of contention. Caramazza and colleagues have interpreted this evidence as indicating that semantic information is "domain-specific" (Caramazza & Shelton 1998; Caramazza & Mahon 2003). Other researchers have adopted different models to explain some of the same findings. Tyler and Moss hold that the selective deficits are an emergent phenomenon. Even though concepts are represented in a unitary distributed system, different types of concepts are structured differently. Since concepts in different domains have different internal structures, impairment of brain function leads to their being differentially affected (Tyler & Moss 2001).

On the face of it, much of this evidence, and the surrounding debate, seems to pertain not to the question of domain specificity but rather to that of brain localization. When damage to a certain part of the brain results in selective impairment in naming animals but not plants or body parts, the question is whether this is evidence that representations underlying our semantic information concerning animals is localized in a particular area of the brain, or whether they are not localized but that some of them are more impaired than others by such damage. Although this is an important question in its own right, it does not bear on domain specificity as such.

Similarly, neuroimaging data that has been brought to bear on this controversy is largely pertinent to the question of localization rather than domain specificity. On the one hand, Caramazza and Mahon (2003, 358) think that "there clearly does seem to be neural differentiation by semantic category" based on neuroimaging data. But Tyler and Moss (2001, 246) find that: "The most striking aspect of the neuroimaging data is the extent to which living and non-living concepts activate common regions with only small and inconsistent differences between domains." The neuroimaging data is obtained mainly by testing healthy

subjects on a variety of tasks (e.g. silent naming, word-picture matching) and then using various techniques (fMRI, PET scans) to determine whether different areas of the brain are differentially involved when processing content derived from different domains (e.g. animals, tools, food items, etc.). But this does not seem to enable us to draw conclusions regarding whether our capacity to think about such domains involves abilities that are generalizable or not. If our knowledge of animals activates different brain areas than our knowledge of tools, that does not mean that any cognitive abilities that range over such domains are restricted to these domains and cannot be applied to others.

Among developmental psychologists, a debate has also raged concerning the domain specificity of our cognitive capacities. For instance, Carey and Spelke have argued that children have innate systems of knowledge that apply to distinct sets of entities and phenomena. Moreover, the domains of human knowledge, such as knowledge of language, physical objects, and number, center on distinct principles. These "core principles" serve to distinguish one domain from another. But despite the fact that Carey and Spelke hold that our cognitive makeup consists of distinct domains, they also claim that conceptual change in these domains occurs in part by constructing mappings between these domains. For instance, mappings between the domains of physics and number play a role in children's reconceptualization of matter and material objects. Though the mapping is slow and difficult, children eventually succeed in using this mapping from one domain to another to differentiate the concept of weight from the concept of density (Carey & Spelke 1994, 191-192). But if one can transplant certain principles of reasoning from one domain to another, then those principles are surely not domain-specific. As I have already argued, it is wrong to argue that a cognitive capacity is domain-specific merely on the grounds that it pertains to a distinct body of knowledge. Rather, generalizability of rules or principles is key, and in this instance that condition would seem to be satisfied, thus casting doubt on whether the capacities in question are truly domain specific (as opposed to innate).

Opponents of the claim of domain specificity also sometimes seem to aim their criticism at a different target. Bates (1994/2001) is at pains to distinguish the claim of domain specificity from claims of innateness and (brain) localization. She rightly stresses that a cognitive capacity can have any two of these features without the third. However, her characterization of domain specificity is vague; with respect to language, Bates (1994/2001, 134) says that the claim of domain specificity is that "localized language abilities are discontinuous from the rest of mind, separate and 'special'…" Moreover, despite her explicit cautionary notes, in presenting the arguments for and against domain specificity, she sometimes argues against innateness or brain localization instead. For instance, she argues against the domain specificity of language on the grounds that the brain systems that support language show an extraordinary degree of neural plasticity (Bates 1994/2001, 139). But that does not have a direct bearing on the issue of whether knowledge of language or the capacity to learn language can be generalized to other domains. She also characterizes the controversy over the domain specificity of language as follows: "Have we evolved new neural tissue, a new region or a special form of computation that deals with language, and language alone?" (Bates 1994, 138) Whether or not there is a brain region that has evolved to deal with language alone concerns innateness and brain localization rather than domain specificity.

A more promising case for testing this account of domain specificity can be drawn from the research on face recognition. Researchers tend to be divided as to whether the human capacity to recognize the faces of conspecifics is a domain-specific capacity, or whether it is a capacity that is acquired as a result of more general cognitive processes, of the type used to acquire expertise in other areas of human cognition. Without trying to rehearse the voluminous evidence involved, I will mention just two findings that are pertinent to the issue of domain specificity. Humans do not develop expertise for recognizing the hands or bodies of conspecifics that is at all comparable to their expertise for recognizing their faces, as measured by accuracy and reaction time (McKone, Kanwisher & Duchaine 2007, 12). Similarly, humans show decrease in accuracy in identifying faces when those faces are inverted but do not show such a decrease in identifying houses in the inversion condition (Yovel & Kanwisher 2004). The capacity to recognize upright faces rapidly and accurately does not seem to generalize to other visual stimuli. Object recognition is a skill that is in principle generalizable to domains beyond faces (e.g. hands, bodies, houses), but it fails to be so generalizable in humans. This is clearly in keeping with the first and second conditions outlined above.

What of the evolutionary clause in the second condition? Though it is not always explicitly mentioned by the researchers who work in this area, I venture that it is at least implicitly assumed. Consider the following scenario. Suppose it were found that humans can indeed generalize their face recognition capacities to encompass the faces of dogs. Proponents of domain specificity might not give up on their claim that this capacity is domain-specific, but rather insist that it is a domain-specific capacity that is specific to the domain of faces in general, or perhaps mammalian faces. Indeed, even if further evidence came to light suggesting that this extends to other objects like the facades of houses, they might continue to posit that it is a domain-specific capacity dedicated to the detection of objects with certain parts in particular configurations. What would rule out such a challenge? As I argued earlier, there are no ready-made domains that would enable us to dismiss it in principle. Rather, it seems natural to say that such hypothetical data would not be evidence of domain specificity (albeit across a broader domain) because of evolutionary considerations. Since it is likely that such a cognitive ability would have arisen to detect faces rather than, say, faces of humans and dogs (given the relative recency of the domestication of

dogs), any extension beyond the domain of human faces is indeed a generalization of this ability, and an indication that it is not truly domain-specific. In fact, this is explicitly acknowledged by proponents of domain specificity in this area of research. McKone, Kanwisher and Duchaine (2007, 12) hold that the domain-specific theory "proposes that a face template has developed through evolutionary processes, reflecting the extreme social importance of faces."

Conclusion

In this article, I have tried to provide an analysis of domain specificity in cognition that enables us to make a theoretically useful distinction between domain-specific and domain-general cognitive systems. Drawing on examples from the literature, both genuine and spurious, I have tried to show that there are two features that make a cognitive system domain-specific. First, the cognitive system must be one that is in principle generalizable. Hence, it cannot be something like a body of information concerning a particular area, but something more like a rule or principle that has wider applicability. Second, it must be a system that the subject cannot apply to genuinely novel cases, where novel cases are ones of a type which this system was not originally evolved to deal with, or that are not within what has been termed the *proper function* of this cognitive system. This second condition is important in that it provides us with a principled way of delimiting domains, since these are not antecedently given. This condition is meant to help to address the question of whether or not a creature can go beyond the cases for which its cognitive capacities were evolved. The distinction between domain specificity and domain generality matters because a central debate in contemporary cognitive science concerns the extent to which our cognitive capacities are domain-specific tools evolved to solve certain problems in the evolutionary environment, or whether they are domain-general problem-solving capacities. A resolution of this disagreement depends on a clear means of demarcating domain-specific from domain-general system. In addition, it has often been claimed that one of the main points of difference between human cognition and that of other animals is its domain-general nature. Again, this debate cannot be properly adjudicated unless we have a principled way of making the distinction.

References

Bates, E. (1994/2001). Modularity, domain specificity and the development of language. In W. Bechtel et. al. (Eds.). *Philosophy and the neurosciences: A reader*, Oxford: Blackwell (first published 1994).

Boyer, P. & Barrett H. C. (2005). Domain specificity and intuitive ontology. In D. Buss (Ed.). *The handbook of evolutionary psychology*, Hoboken, NJ: Wiley.

Caramazza, A. & Mahon, B. Z. (2003). The organization of conceptual knowledge: The evidence from category-specific semantic deficits. *Trends in Cognitive Sciences*, *7:8*, 354-361.

Caramazza, A. & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience, 10:1, 1-34*.

Carey, S. & Spelke E. (1994). Domain-specific knowledge and conceptual change. In L. A. Hirschfeld & S. A. Gelman (Eds.). *Mapping the mind: Domain specificity in cognition and culture*, Cambridge: Cambridge University Press.

Caro, T.M. & Hauser, M.D. (1992). Is there teaching in nonhuman animals? *Quarterly Review of Biology, 67*, 151–174.

Cheney, D.L. & Seyfarth, R.M. (1985). Social and non-social knowledge in vervet monkeys. *Philosophical Transactions of the Royal Society, B308*, 187-201.

Cosmides, L. & Tooby, J. (1994). Origins of domain specificity: The evolution of functional organization. In L. A. Hirschfeld & S. A. Gelman (Eds.). *Mapping the mind: Domain specificity in cognition and culture*, Cambridge: Cambridge University Press.

Csibra, G. (2007). Teachers in the wild. *Trends in Cognitive Sciences, 11:3*, 95-96.

Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Fodor, J. (2000). *The mind doesn't work that way*. Cambridge, MA: MIT Press.

Frank, N. R. & Richardson, T. (2006). Teaching in tandem-running ants. *Nature*, 439, 153.

Hirschfeld, L. A. & Gelman, S. A. (1994). Toward a topography of mind: An introduction to domain specificity. In L. A. Hirschfeld & S. A. Gelman (Eds.). *Mapping the mind: Domain specificity in cognition and culture*, Cambridge: Cambridge University Press.

Khalidi, M.A. (2001). Innateness and domain specificity. *Philosophical Studies, 105*, 191-210.

McKone, E. Kanwisher, N. & Duchaine, B.C. (2007). Can generic expertise explain special processing for faces? *Trends in Cognitive Sciences, 11:1*, 8-15.

Millikan, R. G. (1989). In defense of proper functions. *Philosophy of Science, 56,* 288-302.

Neander, K. (1991). Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science, 58,* 168-184.

Samuels, R. (1998). Evolutionary psychology and the massive modularity thesis. *British Journal for the Philosophy of Science*, 49, 575-602.

Thornton, A. & McAuliffe, K. (2006). Teaching in wild meerkats. *Science, 313,* 227-229.

Tyler, L. K. & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences*, *5:6*, 244-252.

Yovel, G. & Kanwisher, N. (2004). Face perception: domain specific not process specific. *Neuron*, *44*, 889-898.