Does a Honeybee really have the concept of HIVE?

Bernardo Aguilera Dreyse
PhD Student
University of Sheffield

XIII TIF, Barcelona 2011

1- Introduction

The long standing debate about whether animals can think, has been fuelled in recent years by scientific research that provides strong evidence that the cognitive processes of many animal species are computational. Behaviours that at first sight could appear as immediate responses to environmental contingencies, after careful observation and experimental procedures have shown to be governed by structured representations following complex computational algorithms. Examples come from diverse parts of the animal kingdom, including insects, birds and mammals.

These findings have been received with enthusiasm by advocates of a particular view of animal cognition (Carruthers 2006), which embraces a version of the computational theory of mind (CTM) and the massive modularity hypothesis (MMH). In short, this view holds that animals process information about their environment by means of computation, and that their cognitive architecture is mainly modularised into domain specific processors. What concerns me here is to discuss the claim, put forward by proponents of this view, that most animals have concepts[1], and that those concepts can combine forming (propositional) thoughts. In particular, I will deal with the case of honeybees, for two reasons. First, because their behaviour has been extensively studied and there is general consent that they carry out computations, and secondly because it is particularly provocative to suggest that they can think. If it turns out that honeybees have concepts and thoughts, it appears straightforward to suggest that this capacity is widespread in the animal kingdom.

[1] A terminological note: In philosophical usage "concepts" are generally understood as abstract entities, however in psychology the term is used to designate mental representations (Margolis & Laurence 2007). In this paper I will follow the psychological usage, but understanding particular mental representations as concept *tokens*, that instantiate mental representation *types*. Whenever I refer to concept tokens I shall use italics, and when referring to concept types I shall write them in capitals.

In this paper I will assume that CTM gives a plausible account of human concepts and thoughts, and likewise that there are no reasons in principle to reject the idea that cognitive architecture is massively modular. However, I will argue that when these ideas are deployed for the case of honeybee cognition, in the way put forward by Carruthers, they fail to provide convincing grounds for the individuation of concepts. More precisely, I will claim that two main features of conceptual thought, i.e. stable concept individuation and the generality constraint, are not always satisfied. To present my arguments, I will use HIVE as an example of a possible honeybee concept, because it denotes an elemental component of their environment, that should be present in their repertory of concepts if they have any concepts at all.

This paper is structured in the following way. In section 2 I sketch the basic tenets of CTM and MMH, and in section 3 I will explain how CTM gives a plausible account of conceptual thought, addressing some common objections. These two first sections are intended to give a general background about the views that have inspired the proposal about animal cognition that is criticised further on in the paper. Section 4 gives a brief exposition of current research in honeybee behaviour, in order to make clear how it strongly suggest that they are computational systems. In section 5 I explain the claims put forward by authors who interpret this symbolic processing as a form of conceptual thought, and then in sections 6 and 7 I present my arguments against that view. Finally, section 8 gives some final remarks.

2- The Computational Theory of Mind and the Massive Modularity hypothesis

The CTM has been the dominant view about how the mind works over the past four decades, and it has two basic tenets. One is that the mind is a representational system. That means that the mind picks up information about the environment and encodes it as mental representations. This information is made available by perceptual systems, and can be stored in memory for future processing (Sterelny 1990). The second tenet is that mental representations are processed following computational (i.e. algorithmic) steps. That means that these processes are performed in ways only responsive to the formal properties of the representational states, whilst their contents (i.e. what they mean) are preserved along the computational steps without playing any causal role in the process (Haugeland 1981). One of the most influential articulations of the CTM has been by Fodor (1975, 1987). Since he is often quoted by the proponents of animal cognition, I will focus on his account of CTM for the rest of this paper.

One of the main contributions by Fodor to CTM was to make explicit the idea that the mind must have an inner medium of representation that carries out the computations. He also claimed that the properties of compositionality, productivity and systematicity of thought could only be explained if this inner medium has the structure of a language, which has those properties in virtue of its syntax. For that reason, he proposed that the mind has a language of thought (LOT). According to this view, thinking consist in entertaining sentences in LOT. Words in LOT express concepts, and sentences express propositions. LOT is where thoughts and its properties (i.e. syntactic and semantic) primarily reside. Therefore it is not dependent on any natural language, and its basic structure is innate. This idea has important implications for animal cognition, since it states that thinking is not a capacity derived from the possession of a natural language, leaving open the possibility that non-linguistic creatures could think.

The MMH, by its part, is a claim about cognitive architecture. The main idea is that the mind does not work as a single, domain-general system, but has several functionally distinguishable modules, that process domain-specific information and work fast and rather isolated one from another. Initial accounts of cognitive modules restricted their processing to perceptual and motor information. However, proponents of MMH have proposed that mental processes involving thoughts and reasoning are also modular (Cosmides&Tooby 1994; Pinker 1997).

According to MMH, the modular parcellation of cognitive capacities constitutes a extremely common evolutionary feature, that enhanced the adaptability of organisms by permitting them to deal more efficiently with their environments (Carruthers 2006). That explains why the animal mind is supposed to be massively modular. Some empirical evidence has been put forward to defend this claim. To give one example, the navigational capacities of many animals, including rats and birds, have shown to be modular (See Shettleworth 1998 for a review). They have been studied in artificial environments that offer limited kinds of information that can be used by them to orientate. Animals proved to be able not only to use these different environmental clues to navigate, but to deploy them in a way that requires computation, as vector addition or template matching. However, some kinds of information appear to be perceived and used independently, without the capacity to integrate it with other visual clues. All this suggest that they process the various kinds of spatial information by dedicated cognitive modules, that exhibit the hallmarks of domain-specificity, computational processing and isolation.

3- Conceptual thought in CTM: content and individuation of concepts

According to CTM, thoughts are propositions in LOT, and concepts are the elements from which they are constructed. When an agent is thinking, chains of propositions are tokened in her mind, one leading to the other following algorithmic steps that are sensitive to the syntactic properties of LOT. So, for example, an agent could think:

When it's raining, there are no rabbits in the meadow
Now it's raining
So, there are no rabbits in the meadow

Here, the propositions have a syntactic structure the can be recognised by the system (i.e. the brain) as an instance of *modus ponens*, and then processed in a way that the conclusion follows from the premises. The thought can be carried out mechanically, independently of the content of the concepts involved in it. This suggests that thought can be viewed as a purely syntactic procedure.

This idea of mechanised thinking is at the core of CTM. It has many advantages, one of them being how the logical structure of reasoning could be implemented in a digital computer. However, it has the counterintuitive consequence that what-the-thought-is-about is not playing any causal role in the thinking process. In the previous example, we could replace *rabbits* for *foxes*, and the thinking process will still be the same (i.e. an instance of *modus ponens* specifiable purely by syntax), however, of course it is not the same to think about rabbits than to think about foxes. To conclude that there are foxes in the meadow instead of rabbits can make the difference between life and death for an animal. So, how can semantics be ignored?

The response of CTM is that they simply never said that semantics could be ignored, or that it could be reduced to syntax (Horst 2009). Concepts have semantic properties, that determine their contents, but they are fixed from "outside" the domain of thought, by the input and output causal relations that the concept has with the external world. For example, what makes an agent to instantiate the concept of RABBIT is that she has been caused to think about rabbits every time there has been a causal connection between rabbits and her perceptual systems. In other words, it is the interaction of the agent with the word what give the concepts meaning. Then, in the course of thinking, the semantic properties of the concepts are preserved along the computational processes, but not suppressed. This account of conceptual content is usually called "causal theory of content". There is controversy about how to precisely determine content, and several theories available. However, for the purposes of this paper suffice it to say that a causal factors in determining content are dominant among theorists of CTM (Rey 1997).

It is important to note that even though conceptual content is independent from the syntactic structure of thought, this is not (generally) the case of concept individuation. Two concepts may share their contents (i.e. have the same extension), but differ in two further aspects. One of them is the expression in LOT where they are instantiated, what is usually called “mode of presentation”. For example, an agent can think about rabbits both tokening the expression *rabbit* and *long-eared pet*[2] , which constitute two modes of presentation for the same content. A second aspect is the causal role of these expressions. Two concepts can share the same content, though differ in their causal effects on other thoughts and behaviour. So when an agent thinks tokening the LOT expression *rabbit*, it may be lead to think about hunting, but when tokening *long-eared pet* may be caused to think about the veterinary. Both concepts have the same meaning, but differ in their modes of presentation and causal roles.

Of course, this is not deny that both concepts and propositional attitudes can be individuated just in virtue of their contents. But in the context of a belief-desire psychology concepts are always instantiated in propositional attitudes which, *qua* mental states, are individuated according to three aspects: content, mode of presentation and functional role (Cf. Fodor 1990: Ch.6, 1998: Ch.1). So, concept content is individuated by a causal (extensional) theory. But concepts *qua* mental states are also individuated by their vehicles (i.e. LOT tokens) and the kinds of mental processes they cause (i.e. functionally).

This is by no means uncontroversial, but I believe suitable for the purposes of cognitive psychology, at least for three reasons. First, it captures the intentional dimension of thought, since it is not just an external ascription of concepts, but one that respects the “intensional context” of the expressions in LOT. Secondly, it can explain how there can be mental states realistically constructed. In contrast with views that remain neutral about the nature of mental states (e.g. interpretationism), CTM explains how mental content could be instantiated in an internal medium of representation. This allows us to threat the agent as a genuine thinker, with real and casually efficacious mental states, and also helps to solve the “mode of presentation problem” (i.e. explaining how an agent could have two different thoughts about the same thing, by using different LOT expressions).

Finally, a third advantage of this view is that it can serve the purposes of a scientific psychology, by giving an account of how the contents of structured mental states can take part in the mental life of an organism. LOT provides the cognitive vehicles for causally efficacious sequences of thoughts, whilst their functional roles can describe inferential patterns that instantiate principles of rationality. In other words, it allows a scientific

[2] To simplify the exposition I am using English words as expressions of LOT. But LOT, at least according to the CTM, does not correspond to any natural language.

explanation of how cognitive agents behave in virtue of their intentional mental states (i.e. beliefs, desires, etc.). This is particularly important for this essay, since philosophers of cognitive ethology have embraced a similar view for animal psychology (See Allen&Bekoff 2006). They claim that a mentalistic framework like that used by cognitive psychology (i.e. folk-psychology) can be applied to explain animal behaviour, attributing many animals with structured thoughts that interact causally according with rational patterns.

4- Honeybees as computational systems

In this section I will give a brief review of some complex behaviours that have been studied in honeybees, in order to show how plausibly it seems to claim that they are endowed with computational states and processes.

As is well known, honeybees have notable navigational capacities, that make them able to fly from their hives to sources of food and return. Sometimes they rely on landmarks to orientate, whilst they also use dead reckoning (calculate their position by estimating the direction and distance travelled). They exploit the solar azimuth as a directional referent, being able to estimate its position in the sky at different times of the day in order to set and hold a compass course (Collett&Collett 2002). And more surprisingly, they can integrate this information and use it flexibly. For example, in some experiments bees were captured after feeding and carried in a dark box to an unfamiliar releasing point. When released, they initially continued to fly the course they were on when captured, but they soon recognised they were lost, and begun and extensive search until they found a familiar landmark. Then, they were able to fly straight to their hives, in a vector they have never flown before (Menzel et al 2005). This experiments suggests honeybees can represent many features of their environments, and integrate them with stored information about distance and direction relative to their hives.

Another remarkable fact about honeybees are their communicational capacities. Foraging bees transmit information about food resources to other bees via different kinds of dances they perform inside the hive (Gould&Gould 1996). Some features of the dance such as the angle of movement, as measured from the vertical, and the number of “waggles” they made at some point of the dance, convey information about the expected direction of the sun for the time of the day, and the distance of the food source. And the bees in the hive are not just able to integrate the communicated information and fly to the food, but can also evaluate it along a number of dimensions. For example, they are less likely to fly to distant sources of food, and show preference for rich sources of food.

This findings suggest that the behaviour of honeybees cannot rely exclusively on fixed action patterns, or be conditioned responses to stimuli. Instead, they can form complex and structured representations of their environments, including information about distance, time, direction and location. They can also transmit this information and use in a rather flexible and systematic way. Plausibly, many authors have claimed that the best explanation for these complex behaviours is that honeybees can carry out computational processes over causally efficacious and structured representations (Carruthers 2006; Gallistel 2009; Tetzlaff&Rey 2009).

5- Honeybees as thinking creatures

Through several writings, Carruthers (2004, 2006, 2009) has given a detailed defence of the computational capacities and the massive modularity of animal cognition. Among them are honeybees, whose striking behavioural complexities I summarised above. He also moves a step forward claiming that honeybees have conceptual thought, according with the framework of CTM. I will summarise his view, to then present my arguments against it in the following sections.

He argues that the capacity of certain animals to represent specific features of their environment and to process them following algorithmic steps, constitutes a genuine form of means/ends reasoning. He contrasts it with forms of associative conditioning or innate releasing mechanisms, which cannot explain the flexibility and complexity of certain behaviours (e.g. those of the honeybees presented above). On the contrary, many animal behaviours are mediated by cognitive processes that involve explicit representations and purposeful reasoning. He goes on to claim that these processes can be characterised in terms of belief states and desire states that are discrete, structured, and causally efficacious in virtue of their structural properties.

The force of Carruthers' argument rests on the assumption that CTM works for human beings. Since there are no problems in constructing a belief/desire psychology based on computational processes carried out in LOT, the same could be done for the animal case. The difference between humans and other animals, he claims, is basically matter of complexity and not of kind. The grounds for having a LOT with syntactic and semantic properties are not defined by the number or complexity of the representations and computations involved. It suffices to have a simple compositional medium of representation, able to structure belief and desire states that can interact through basic practical inferences to select and guide behaviour.

As previously noted, Carruthers' idea of animal cognition also involves claims about massive modular architecture. The computational systems of animals are supposed to be organised in cognitive modules, and that implies that the representations and computations the animal carries out are distributed into separated, domain-specific and rather isolated units. For example, honeybees seem to use distinct modular systems to navigate inside or outside their hives. When inside the hive they orientate themselves using gravity-based and olfactory cues, whilst they rely on solar bearings when outside it. And it is not that they *choose* between one navigational system or another. According to MMH, honeybees do not represent those systems as alternative sets of spatial representations[3], but they activate one or another when the relevant input is present. So, if honeybees are thinkers, they deploy different ways of thinking about their environment depending on which module they are using. Those differences are principally related with the representational vehicles they deploy (e.g. from distinct perceptual formats), and the functional roles they occupy (i.e. directed to different domain-specific tasks).

6- First argument against Carruthers: stable concept individuation

Now I turn to argue that the view that honeybees have conceptual thought, as explained in the preceding section, has several problems. In particular, I will argue that it becomes implausible to state that honeybees can individuate and use the concept of HIVE. Since this concept denotes an essential feature of their environment, it is hard to see how they could be treated as genuine thinkers if they lack it.

My arguments appeal to problems that a massively modular architecture presents for the individuation of a concept. Recall the two navigation modules that honeybees have for orientating inside or outside the hive. They should be able to entertain the concept of HIVE in both cases, presumably in different representational formats, one based on gravitational and olfactory cues, whilst the other based on visual cues. Besides, they should be able to combine this concept with others, in order to form propositional thoughts. According with the framework of MMH, instances of the concept of HIVE should coexist within each module with other concepts that concern the specific computations that the module was designed to perform. For example, suppose that the navigational module for outside the hive has the concept of BLUE, whilst the module for inside does not. Instead the module for navigating inside the hive has the concept of WAX-ODOUR, absent in the other module. This observation leads to the conclusion that the honeybee would only be able to think *the hive has wax-odour* when it is inside the hive, whilst the thought *the hive is blue* could only be entertained when outside it.

[3] Because they lack second order beliefs. In Carruthers (2002), the author defends his view that cross-modular cognitive integration of thoughts is restricted to linguistic creatures.

Does that mean that the honeybee has two different concepts of HIVE, one for each module? If we recall how CTM defined the semantic properties of LOT, we could give an initial response: they are instances of the same concept, since they both have the same extension, i.e. they are both about hives. However, their mode of presentation and functional role in the propositions they constitute must be different, since each module has a domain-specific set of representations and functional roles. They respond to specific input and output channels, and carry out the computations specified by the function of the module. This could be seen as a problem when individuating the concept according with its psychological significance.

However, Carruthers sees this situation as unproblematic, suggesting that animals do not have a single LOT, but several LOTs, one for each module. That means that the honeybee may turn out to have many modules where can think about hives, but does it in radically different ways, as if there were different languages that cannot understand each other (even in the case of being provided with a communication channel). However, I believe that this idea does not work. The idea of multiple modules processing the concept of HIVE in the way proposed by Carruthers makes the individuation of this concept problematic. My argument can be summarised as follows:

There are several modules that have the concept of HIVE within a single honeybee
Each concept of HIVE has its own causal role and expression in LOT
Concepts are partly individuated by causal role and expression in LOT

Therefore, a single honeybee have several concepts of HIVE

What fixes the content of HIVE is its extension. Two agents share the content of HIVE if hives in the world causally co-vary with the instances of HIVE in their heads. But, as I explained in section 3, concepts, as they figure in mental states, are not individuated just by their contents. Content attribution depends on their extension, whilst concept attribution depends on their place in LOT and cognitive background. More precisely, concepts are relations between a cognitive agent with: a LOT expression, its content, and functional role. Of course that does not lead to the problematic consequences of meaning holism. Two agents can have different concepts of *hive*, and a single agent may change her concept of *hive* from one day or another, but their meanings will remain the same because their truth-conditions for propositions where HIVE participates will still be equivalent. However, what becomes problematic is that a single cognitive agent has two (or more) concepts of *hive*, as would happen in the case of the honeybee. Even if they share the same content, the concepts may have conflicting functional roles and therefore cannot be individuated as the concept of HIVE.

So my point here is that we cannot account for a stable individuation of a concept in such a fragmentary way. According to the picture given by the CTM, instead, LOT provides cognitive agents with the medium for representing and thinking about the world. Even if we can think about the same thing in two different ways (recall the example of RABBIT given in section 3), and thus individuate the same content in two different concepts, LOT ensures that they conform two perspectives or points of view about the same object. I believe that the very ideas of mode of presentation and functional role presume a common LOT where representations and roles are instantiated. In the case of the honeybee, on the contrary, vehicles and functions are enclosed in each module, and so the same counts for concept tokens. It strikes me as highly implausible that the same concept type (e.g. HIVE) could be multiple individuated (e.g. *hive1*, *hive2*, *hive3*, etc.) with incompatible vehicles and functions and still regard the agent as genuinely capable of entertaining the concept of HIVE. The present argument is closely related with the one presented in the next section, and I hope will become clearer in this context.

7- Second argument against Carruthers: the Generality Constraint

My second argument is also related with some consequences of MMH in the individuation of concepts, but this time I focus on the generality constraint (GC). The GC is often assumed as an essential characteristic of conceptual thought, and was first stated by Evans (1982) as follows:

> We cannot avoid thinking of a thought about an individual object x, to the effect that it is F, as the exercise of two separable capacities; one being the capacity to think of x, which could be equally exercised in thoughts about x to the effect that it is G or H; and the other being a conception of what it is to be F, which could be equally exercised in thoughts about other individuals, to the effect that they are F. (p. 75)

The main idea is that genuine thinkers should be capable to produce and entertain an unbounded set of novel well-formed combinations of concepts. This capacity is closely related with what has been called the systematicity and productivity of thought, which have been proclaimed by CTM theorists as elemental features of thought. In Fodor's words:

> Productivity and systematicity are also universal features of human thought (and, for all I know, of the thoughts of many infra-human creatures). There us no upper bound to the number of thoughts that a person can think. (1994: 106-7)

Moreover, CTM offers one of the most compelling explanations about the cognitive mechanisms that underlie this features, based on the compositional nature of LOT (see section 3). So the GC can be safely regarded as a hallmark of thought that honeybees should fulfil if they have genuine concepts.

Now suppose that the perceptual apparatus of the honeybee is able to discriminate between three colours: green, yellow and red, and that this capacity is deployed in a module for flower recognition. Also suppose that the navigational module for outside the hive includes among its domain-specific repertory of representations just the colours *green* and *yellow*, but not *red*. So, among the operations of this module the honeybee might be able to combine the concept of HIVE with the concepts of GREEN and YELLOW, forming the thoughts *the hive is green* and *the hive is yellow*. However, she will not be able to think *the hive is red*. This appears to violate the GC.

Carruthers (2004, 2009) has defended the conceptual capacities of animals, claiming that it is perfectly possible for many of them to satisfy the GC, since they are capable of forming thoughts with compositional structure. The fact that they do not ever form certain new relations between concepts may be just a consequence of them not being *interested* in those relationships. But following a “weak” version of the generality constraint, he suggests that it be *metaphysically* possible (MP) for honeybees to combine their concepts in new ways, and thus realise the GC. Note that the requirement that this capacity must be *metaphysically* possible is stronger than the claim that it must be *logically* or *conceptually* possible. Carruthers dismiss the second requirement because it is uninteresting whether it is possible for honeybees to realise this capacity through acquiring new powers of representation (e.g. metarepresentations). His claim is that it is MP for honeybees, given some variations in their actual cognitive architecture, to realise the GC. So, even if it is actually impossible (concerning their cognitive architecture) for honeybees to think *the hive is read*, given the appropriate circumstances they may well be able to entertain this thought.

I find this argument unconvincing. My reply goes like this:

The GC could be satisfied if new cross-modular conceptual combinations are MP
New cross-modular conceptual combinations are MP iff concepts could be detached from their current role
It is not MP that concepts be detached from its current role in their respective modules

---

Therefore, the GC cannot be satisfied

The main assumption of my argument is that concepts cannot be detached from their current cognitive domain, that is, from the functional role they have in their respective modules. The reason for that is that even if cross-modular concepts share the same extension, they can have radically different representational formats and functional roles.

In other words, they were *built* for different purposes, and so it is not possible, even in principle, to combine with other concepts outside their proper domain.

Even the metaphor of people speaking different languages in each module may be not accurate, since in that case, with the help of a dictionary, they could translate and understand each other. Rather, the picture of cognitive modules suggested by MMH is of processing units that are structurally different, with specialised machinery for their own particular domains. A common analogy is with the body (Pinker 1997), which is organised into distinct organs, assembled by natural selection for particular tasks. So, for example, suppose that under a functional description the liver is able to represent certain molecules that are in the blood and process them. An analogue case can be a kidney, where molecules in the blood are also sensed and processed. The idea of combining representations from different modules would be equivalent to combine representations the kidney has of certain molecules with representations from the liver. They are specified in terms of specific machinery of the organs (i.e. receptors, enzymes, etc.), encoded in representational formats that are sensitive to particular features of molecules in virtue of the role they play in the processes the organs were *built* to perform. Therefore, it is hard to imagine how their representations could be put together to compose new representations. The only way to carry out this idea could be to deploy an external representational medium where the representations of both organs figure and combine. But this implies that it is not MP for the organs, by using their actual representational systems, to combine their representations.

8- Final remarks

As I have stated since the beginning, my purpose in this paper has not been to criticise the main tenets of CTM or MMH. They could be perfectly true, and some version of them suitable for animal cognition. My point has been to argue that the requirements for conceptual thought are not fulfilled in a model that simply conjoins both views without further refinements.

It is always tempting to attribute a belief-desire psychology to animals, and without doubt the evidence about their computational capacities makes them good candidates to be thinkers. However, whether this evidence alone is enough to account for the conceptual nature of their representations is far from clear. At least, it has been the purpose of this paper to show that Carruthers' account of honeybee concepts presents at least two problems.

Probably, one of the main limitations of MMH when accounting for minded organisms is that it contradicts the rather intuitive idea that minded organisms have an autonomous, unitary perspective of their world (Cf. Crane 2001). If one is to buy Carruthers' picture of a mind where concepts and thoughts are realised in multiple modules, then the animal mind consist in a collection of input-output independent processors, and it becomes hard to see how from this partitioned system animals could develop a meaningful point of view of their surroundings. As I mentioned in section 6, the very ideas of mode of presentation and functional role seem to presuppose a common representational medium where alternative modes and functions are instantiated.

A plausible alternative, I believe, could still hold that animal cognition is massively modular, but claim that genuine minds emerged when second-order representations (or metarepresentations) evolved in animals. This could have provided a cross-modular medium to detach the split repertoire of representations contained in modules, and integrate them into a unitary representational system, that gets closer to an intuitive picture of what a mind is. Some authors have suggested that metarepresentational capacities are present in just a few highly intelligent animals, such as some primates (Sperber 2000). However, others have proposed that metarepresentations could be wide spread in the animal kingdom, probably under a non-propositional representational format (Bermúdez 2009; Proust 2009). The issue about whether metarepresentations are a necessary condition for having a mind (and therefore genuine concepts) or when they appeared in phylogeny, goes beyond the purposes of this paper, however.

9- References

Allen, C., & Bekoff, M. (1997). Species of Mind: The Philosophy and Biology of Cognitive Ethology. Cambridge, Massachusetts: MIT Press.

Bermúdez, J. L. (2009). Mind reading in the animal kingdom. In R. Lurz, *The Philosophy of Animal Minds* (pp. 145-164). Cambridge: Cambridge University Press.

Carruthers, P. (2002). The cognitive functions of language. *The Behavioral and brain sciences*, *25*(6), 657-74; discussion 674-725.

Carruthers, P. (2004). On Being Simple Minded. *American Philosophical Quaterly*, *41*(3), 205-220.

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

Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz, *The Philosophy of Animal Minds* (pp. 89-107). Cambridge: Cambridge University Press.

Collett, T. S., & Collett, M. (2002). Memory use in insect visual navigation. *Nature reviews. Neuroscience*, *3*(7), 542-52.

Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: the evolution of functional organization. *Mapping the mind: Domain specificity in cognition and culture*, 85–116.

Crane, T. (2001). *Elements of Mind*. New York: Oxford University Press.

Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.

Fodor, J. (1975). *The Language of Thought*. Cambridge, Massachusetts: Harvard University Press.

Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. London: Bradford.

Fodor, J. (1992). *A Theory of Content and Other Essays*. Cambridge Mass: MIT Press.

Fodor, J. (1994). Concepts: a potboiler. *Cognition*, *50*(1-3), 95-113.

Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.

Gallistel, C. (2009). The foundational abstractions. In M. Piattelli-Palmirini, J. Uriagereka, & P. Salaburu, *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country* (pp. 58-73). New York: Oxford University Press.

Gould, J., & Gould, C. (1998). *The Honey Bee*. (Scientific American Library). New York.

Haugeland, J. (1979). Understanding Natural Language. *Journal of Philosophy*, *76*, 619-632.

Horst, S. (2009). The Computational Theory of Mind. In E. Zalta, *The Stanford Encyclopedia of Philosophy*. Retrieved from http://plato.stanford.edu/archives/win2009/entries/computational-mind.

Margolis, E., & Laurence, S. (2007). The Ontology of Concepts—Abstract Objects or Mental Representations? *Noûs*, *41*(4), 561-593.

Pinker, S. (1997). *How the Mind Works*. New York: Norton.

Proust, J. (2009). The representational basis of brute metacognition: a proposal. In R. Lurz, *The Philosophy of Animal Minds* (pp. 165-183). Cambridge: Cambridge University Press.

Rey, G. (1997). *Contemporary Philosophy of Mind: A Contentiously Classical Approach*. Oxford: Blackwell.

Shettleworth, S. (1998). Cognition, Evolution, and Behavior. New York: Oxford University Press.

Sperber, D. (2000). Metarepresentations in an Evolutionary Perspective. In D. Sperber, *Metarepresentations: A Multidisciplinary Perspective* (pp. 117-138). New York: Oxford University Press.

Sterelny, K. (1990). *The Representational Theory of Mind: An Introduction*. Oxford: Basil Blackwell.

Tetzlaff, M., & Rey, G. (2009). Systematicity and intentional realism in honeybee navigation. In R. Lurz, *The Philosophy of Animal Minds* (pp. 72-88). Cambridge: Cambridge University Press.