DO CONCEPTS EXIST?
A NATURALISTIC POINT OF VIEW

MARCELLO FRIXIONE

Introduction

In this article, I argue that: the idea that within the mind there exists a collection of more or less homogenous structures of some kind, corresponding to our commonsense notion of "concept", is likely upon investigation to prove wrong.

In the literature in the field, the term *concept* is used to denote kinds of mental representation that *prima facie* seem very different. The fact that these phenomena sound intuitively heterogeneous would not be cause for any problem if some underlying cognitive mechanism were discovered. However, if we try to either a) identify some common cognitive mechanism or mental structure underlying such phenomena, or b) specify a more restrictive and homogeneous notion of "concept", the results are rather poor.

The research on concepts is generally considered to be a successful example of the interdisciplinary approach of cognitive science. However, when we try to single out which results have been achieved, we realise that the main (perhaps the only) ascertained fact is the inadequacy of the classical approach to categorisation. The failure of the classical approach has not involved the emergence of a clear alternative point of view. Even if the existence of prototypical effects is an ascertained fact, many psychological data hardly seem compatible with the hypothesis of a uniform underlying cognitive mechanism.

In my opinion, this is due to the fact that prototypical effects, far from being a symptom of some homogeneous cognitive structure, are more plausibly the effect of some "ecological constraints" on the mind. In other words, our way of categorising the world does not exhibit prototypical effects because our mind is structured in a particular way; rather,

prototypical effects are determined by the fact that the world interacts with our mind in a certain way.

1. The state of the art

In recent cognitive science literature, one happens to read bold statements like the following:

> Concepts are the most fundamental constructs in theories of mind. (Laurence and Margolis, 1999: 3).

> Concepts are the glue that holds our mental world together (Murphy, 2002: 1).

> [Concepts are] The elements from which propositional thought is constructed, thus providing a means of understanding the world (Hampton, 1999: 176)

In spite of similar enthusiastic claims, I suspect that from the standpoint of a natural science of the mind (i.e., from the standpoint of cognitive science), the notion of *concept* could turn out to be spurious. The idea that there exists within the mind a collection of more or less homogenous structures exists, corresponding to our commonsense notion of "concept" is likely to turn out to be wrong. In other words, I cast doubt on the thesis that concepts form a natural kind from the standpoint of cognitive science.

The study of concepts is often quoted as a paradigmatic model of interdisciplinary research in cognitive science, involving contributions from psychology, philosophy, anthropology, linguistics and artificial intelligence. Moreover, it is generally considered to be a successful example of such an approach. However, if we try to single out the results achieved, we realise that the main (perhaps the only) ascertained fact is how inadequate the classical approach to categorisation really is: the great majority of researchers agree that most concepts used by humans in everyday practice cannot be characterized in terms of necessary and sufficient conditions. There is near-unanimous agreement that ordinary and classical categories differ in some crucial respect: most everyday categories do not have clear-cut boundaries; they exhibit a graded structure in the sense that some entities are better specimens of the category than others, and so on[1]. Besides these (mainly negative)

[1] For positions that deny the relevance of graded structures and prototypical effects in characterising concepts, see Fodor and Lepore (1996) and Conolly *et al.* (2007).

achievements, the failure of the classical conception did not involve the emergence of a clear alternative point of view.

In the technical literature, the term *concept* is used to denote phenomena that *prima facie* are rather different. It so happens that among *concepts* are counted:

- any mental representation of classes of (natural or artificial) physical entities, of actions or events, of abstract entities, of social institutions, and so on;
- mental representations of abstract notions, such as *logical concepts* (disjunction, negation, …), or notions like *cause*, *truth* (sic!), and so on;
- lexical concepts, i.e., the mental representations of the meanings of the words of natural languages;
- mental representations associated to arbitrary (definite or indefinite) descriptions, such as *ways to escape being killed by the Mafia*, *my best friend*, *food that people usually eat at Christmas*, and so on;
- representations of sets of arbitrary stimuli, such as patterns of dots, strings of digits and letters, geometric figures (many experiments on categorisation make use of this kind of stimuli).

Of course, the entries in the list above are not mutually exclusive. Logical concepts and the mental representations of many classes of entities are in their turn lexical concepts. Classes of physical or abstract entities can be represented using descriptions, and so on. However, the list sounds intuitively heterogeneous. This would not be cause for any problem if some underlying cognitive mechanism were discovered. But if we try to either a) identify some common cognitive mechanism or mental structure underlying such phenomena, or, b) specify in cognitive terms a more restrictive and homogeneous notion of concept, the results are rather poor.

Significant from this point of view are the conclusions of Murphy (2002), a book that in many respects reflects the current status of empirical research on conceptualisation. Murphy takes into account the three main approaches to the study of concepts, namely the *prototype view*, the *exemplar view* and the *knowledge approach*.

According to the *prototype view*, categories are mentally represented in terms of *protoypes*, i.e. of explicit representations of the "best exemplar" of the category. It is often assumed that prototypes are mentally coded as lists of features. For example, the prototype of the category of birds could

be a list including such features as *body covered with feathers*, *flies*, *lays eggs*, and so on.

According to the *exemplar view*, the mental representation of a concept consists of a set of representations of exemplars that are stored in memory. For example, according to this approach, my representation of the concept *dog* corresponds to the set of mental representations of the dogs that I have encountered in my lifetime. Of course, we can suppose that not all the dogs that I remember carry the same cognitive weight, and that the representation of the various instances can be more or less complete.

According to the *theory-theory view* (also known as the *knowledge approach*) concepts are part of our general theories of the world. Concepts are defined by their place in the general system of our beliefs: every time we acquire some information concerning a given concept, this involves some change in the concept itself.

Researchers carefully distinguish between *prototypes* and *prototypical effects*. The existence of prototypical effects is an ascertained fact strictly related to the failure of the classical approach, while the existence of some set of homogeneous mental constructs called *prototypes* is only one of the possible hypotheses for explaining such effects. All the above-mentioned views in some way account for prototypical effects, while only the prototype view endorses the thesis that prototypes exist as specific mental constructs.

In his book, Murphy contrasts the above approaches in relation to different classes of problems, including learning, induction, lexical concepts, children's concepts, and so on. His conclusions are rather discouraging: the result of comparing the various approaches is that "there is no clear, dominant winner" (ibid., p. 488).

Murphy's reaction to this state of affairs is ecumenical. He proposes combining the prototype view and knowledge approach: "our theory of concepts must be primarily prototype-based. That is, it must be a description of an entire concept, with its typicality features (presumably weighted by their importance). However, this description must be part of a broader knowledge representation scheme in which the concept is positioned both within a hierarchy (subordinates, superordinates, and coordinates) and within a theoretical framework (or frameworks) appropriate to that domain" (ibid., p. 488). Moreover, "exemplar models can account for certain category learning-results better than prototype models can" (ibid., p.490).

Therefore, Murphy proposes a mixed theory of concepts, which combines aspects of the prototype view, the exemplar view and the knowledge approach: "On the proposal I am making, people attempt to

form prototypes as part of a larger knowledge structure when they learn concepts. But at the same time they remember exemplars and these memories may influence them in a variety of ways" (ibid., p. 492).

The solution favoured by Murphy is the following: "in order to emphasize the relation between knowledge and prototypes, I would suggest that *schema representation* is more appropriate than the simple list of features that is often used" (p. 489, my italics). However, the notion of schema offered by Murphy is extremely vague; schemata are characterised in a way that is very similar to the frames proposed by Marvin Minsky (1975) and adopted with poor results in many AI projects in the seventies. Thirty years on, it seems unlikely that this direction could turn out to be promising.

Summing up, the situation is rather demoralising (in Murphy's words, "concepts are a mess", ibid., p. 492). Further problems separately concern the various approaches that concur with Murphy's hybrid solution. As an example, consider the knowledge approach. Either we can find some *empirical* criterion that allows us to distinguish conceptual knowledge from non-conceptual (say, factual) information, or we are in serious trouble. Indeed, unless we try to revive the analytic-synthetic distinction, we have to embrace some form of holistic position, and in my opinion holism is rather incompatible with a realistic attitude towards concepts (realism towards concepts seems to involve our being able to single out concepts from non-conceptual knowledge).

Moreover, many psychological data seem hardly compatible with the hypothesis of a uniform mental representation of categories. For example, prototypical effects showed a great sensibility to context (Barsalou 1987), and in different domains prototypical judgments seem to be due to different mechanisms (as an example we have the so-called goal-derived categories studied by Barsalou 1985).

2. Prototypical effects and ecological constraints

My claim is that the above-mentioned difficulties are due to the fact that prototypical effects, far from being a symptom of some homogeneous cognitive structure underlying categorisation, are more plausibly the effect of "ecological constraints" acting on the mind. In other words: our way of categorising the world does not exhibit prototypical effects because our mind is structured in a particular way; rather, prototypical effects are caused by the fact that the world interacts with our mind in a certain way. Prototypical effects emerge as a result of the constraints that the environment imposes on any (finite) agent interacting with it. As such,

they do not reveal any particular, specific feature of the human mind: all cognitive agents that have to survive in their environment would develop categories with a similar structure in order to solve similar problems. Obviously, this does not say very much about *how* these problems have been solved by the human cognitive system. Nor does it involve some common cognitive mechanism underlying prototypical effects.

In order to explain the "ecological" character of the constraints at the base of prototypical behaviour, let us consider the following cases, which are in some sense paradigmatic. In order to accomplish their purposes, cognitive agents may classify their environment according to categories that:

a) are intrinsically fuzzy (or vague); or
b) do not correspond to "natural kinds"; or
c) correspond to natural kinds whose essence is (totally or partially) unknown.

Case a) is the less problematic. It includes categories that establish arbitrary discontinuities in phenomena that are intrinsically continuous. A paradigmatic example is colour categories: the chromatic spectrum is a continuum, and any way to categorise colours would give rise to vague concepts. Such categories cannot have clear-cut boundaries. Consequently, there must be "better" and "worse" exemplars, and therefore forms of prototypical effect arise.

Point b) includes cases in which prototypical effects are determined by the fact that categories do not "capture essences", do not "carve the world at its joints". In order to satisfy their needs, cognitive agents are often interested in grouping together entities in ways that do not correspond to the "deep structure" of the world. Paradigmatic examples of b) are artefacts, tools, various types of actions and events, and many social categories. Case b) also includes categories whose members are natural entities which, however, are grouped together according to criteria that are extrinsic to the organisation of the natural world. Good examples of this kind are food categories. Consider such terms as:

- *légumes* (French), *verdura* (Italian), *vegetable* (English)
- *fruit* (as opposed to *vegetable*)
- *pescado* (Spanish)
- *frutti di mare* (Italian), *mariscos* (Spanish), *seafood* (English)
- *viande* (French), *carne* (Italian)

Légumes, for example, include quite heterogeneous parts of vegetal organisms such as seeds, fruits, leaves, flowers, roots and branches. In the category of *mariscos* (or *frutti di mare*) almost all animal phyla are represented: mainly *Mollusca* and *Artropoda* (*Crustacea* above all), but also *Chordata*, *Echinodermata* and *Coelenterata*.

It could be argued that, even if such categories do not correspond to natural kinds from a strictly biological point of view, they could rather be considered "substances" from a "biochemical" standpoint. *Viande*, for example, could be any animal tissue (of mammals or birds) that, by virtue of its chemical features, can be used as food by humans. However, this does not constitute a sufficient condition. In order to consider some matter as *viande* (or *légumes*, or *seafood*, and so on) further constraints must be satisfied. Its taste must not be disgusting, it must not be dangerous due to the presence of poisons or parasites, it must be easy enough to collect in satisfactory quantities, and so on (leaving apart cultural and anthropological constraints, maybe partially motivated by practical considerations, such as the rules of *kasherut* for Jews or the prohibition of cannibalism – cfr. Harris 1985). For example, *Amanita caesarea* is an – indeed rather atypical – example of *verdura* (it is not even a plant), but the poisonous *Amanita phalloides* cannot be considered *verdura* at all, in spite of its biological similarity to the former. Similarly, *Littorina neritoides* is not *seafood* because of its small size, while *Littorina littorea* counts as a good example of this category, in spite of the extreme biological, ecological and taxononomical similarity of the two species.

Case c) is that of natural kinds whose essence is (totally or partially) unknown to a cognitive agent; in such cases, the agent uses stereotypes (as per Putnam 1975) that are partial and fallible, and which therefore give rise to uncertain and dubious cases. Cognitive agents that for some reason are interested in categorising, say, iron, but do not know that iron is the element with atomic number 26, adopt fallible identification procedures that work well in some cases (the "typical" ones) and poorly in other cases. This is related also to the fact that agents are usually more interested in some forms of a natural kind than others (e.g, liquid, drinkable water vs. steam).

From the cognitive point of view, c) is not dissimilar to the case of categories with a classical, clear cut definition, which however is unknown by some cognitive agent. Imagine an undergraduate student who knows that some members of the faculty are full professors, without exactly knowing how *full professor* is defined. Our student would probably develop a "stereotype" that will help deal with the academic world: full professors are usually older than their colleagues, they have greater

authority and power, sometimes they are more erudite and have greater scientific prestige, frequently they are more arrogant, and so on.

In both b) and c) prototypical effects do not depend (only) on the vagueness of the employed criteria, but also on the "bad correspondence" between identification criteria and the word. Criteria are intrinsically imprecise, and give rise to dubious cases. In any case, all the above-mentioned cases a), b) and c) give rise to prototypical effects.

Another aspect of cognition in some way related to prototypical effects in categorisation lies in inferences in which subjects "jump to conclusions" on the basis of default information. For example, birds, by default, are able to fly; therefore, given that Tweety is a bird, I jump to the provisional conclusion that Tweety is able to fly. Of course, I must be prepared to withdraw this conclusion should Tweety turn out to be atypical in his flying abilities (because, for example, he is an ostrich, or a penguin, or a featherless chick). Such an inference from incomplete premises is an example of what is called non-monotonic reasoning (see e.g. Koons 2005).

Default values that allow the drawing of non-monotonic inferences are strictly akin to prototypical knowledge. In other words, associating prototypical information to concepts turns out to be useful in order to make defeasible inferences starting from incomplete knowledge. Default reasoning is a widespread cognitive phenomenon that is, in its turn, determined by ecological constraints, namely by the need for cognitive agents to face various kinds of problems with limited access to relevant knowledge.

Default information is involved in countless cognitive tasks, such as, for example, reasoning on actions and events, and the frame problem, or the default assumptions built into the visual system in order to solve the "hill posed" problem of reconstructing the properties of distal stimuli starting from incomplete proximal information. It is likely that default "inferences" are exploited at many different levels of the cognitive architecture in completely independent components.

On the basis of the above considerations, we can conclude that not only do prototypical effects not say much about the cognitive structures that determine them, but a given cognitive system may face similar problems in different ways, and that, within the same system, prototypical behaviours in different cases are determined by different mechanisms.

It is well known from psychological experiments that human behaviour exhibits prototypical effects also in the case of fully classical categories (for example, categories from arithmetic or geometry). This could be used against the thesis that prototypical effects depend on ecological constraints

and do not correspond to any specific feature of the cognitive representation of categories: the fact that people exhibit prototypical behaviours even for classical categories could be explained in terms of some deep, unitary feature of human concept representation. However, this is not necessarily the case. We can account for prototypical effects for classical categories adopting a distinction similar to the core/identification procedures pair proposed by Miller and Johnson-Laird (1976)[2]. There are cases in which the classical definition of some category (its core) is well known, but in many practical cases it cannot be applied, and people must trust to fallible criteria. For example, even if I know perfectly well what a *grandmother* is, when I categorise somebody as a grandmother I can seldom use the definition, and I must often rely on fallible criteria: grandmothers are elderly people with white hair, they are kind to children, and so on. In some sense, this example is similar to case c) above. Moreover, it is likely that also in the case of formal (mathematical) concepts prototypical effects depend on identification procedures. Let us consider even numbers. We tend to consider number 4 as a more "typical" specimen of this category than, say, 351274195336. This is probably due to the fact that small even numbers (2, 4, 6, 8…) are presumably explicitly stored in our memory in some sort of lookup table, and are therefore easy to check and generate. For larger numbers (as 351274195336) we need a procedure ("check the last figure on the right") that, albeit simple, is more demanding (at least because it requires the procedure for small numbers as a part of the whole process). Compare this with the case of prime numbers: small prime numbers (2, 3, 5, 7, 11…) are explicitly stored in memory, and we can recognise them at a glance, whilst for larger numbers (e.g., 2750159) a rather complicated and time-consuming procedure is needed. The case of geometrical figures is analogous. Consider figures i) and ii) below.

[2] On this point, see for example Armstrong *et al.* (1983) and Rey (1983).

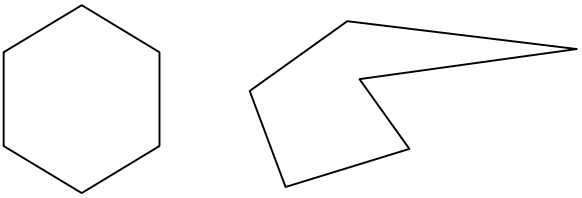

i) ii)

We can immediately recognise without any effort i) as a hexagon (maybe because it has a "good Gestalt", or maybe because we have been educated at school to take into account convex, preferably regular polygons). In order to decide that ii) is a hexagon, we must explicitly count the sides in a sequential way.

That prototypical effects do not correspond to some specific feature of mental organisation is witnessed also by the fact that, in some categorisation experiments, people dealing with abstract stimuli carefully prepared by experimenters to favour the emergence of prototypical effects, prefer to individuate classical categories (Medin *et al.* 1987). (In my opinion, these results show that our mind has no particular bias against classical categories. Rather, the problem is that classical categories are not particularly useful, in that they do not fit well with the world – at least as far as we can access it: in most cases, classical categories are completely ineffective in relation to our needs.)

3. Do concepts exist (as a natural kind)?

In my opinion, that prototypical behaviours in different cases could be determined by different mechanisms is not a mere theoretical possibility. It is what is likely to happen in actual cognitive systems (namely, in the *human* cognitive system), as is suggested by many empirical results of psychological research on concepts. I have already mentioned the goal-derived categories studied by Barsalou (1985). Goal-derived categories are characterised in terms of how their members accomplish some given goal; possible examples are *things to eat on a diet* or *ways to escape being killed by the Mafia*. Barsalou stresses that such categories can hardly be explained in terms of some notion of similarity based on lists of perceptual features, as is usually the case in prototypical explanations of categories grouping together physical objects or artificial stimuli.

The hypothesis that prototypical effects could not correspond to some specific mental structure or to some specific feature of the cognitive system is corroborated by the observation that in computational artificial systems some form of prototypical effect can be obtained using very different computational devices. Forms of "prototypical representation" can be obtained adopting, for example, some non monotonic logical formalism, symbolic models of associative memory such as frames or inheritance networks, artificial (supervised or unsupervised) neural nets, fuzzy systems, probabilistic representations, and so on.

In section 1 I observed that concepts in cognitive science form a *prima facie* rather heterogeneous group. It is likely that some "concepts" fit well with the assumptions of the knowledge approach, in that they (at least partially) depend on global knowledge of the system. As observed before, in such cases it is unlikely that "conceptual knowledge" can be singled out from factual knowledge.

Fortunately, probably not all "concepts" are holistic in nature. Parts of our categorisation abilities seem to depend on factors that can be circumscribed to particular structures of the cognitive architecture. Nevertheless, such structures are likely to depend on different aspects of cognition.

Some "concepts" have a rather specific perceptual counterpart. In addition, in some cases this perceptual counterpart depends on factors that are both innate and shaped by precise anatomical or cognitive factors. As an example, consider chromatic categories: basic level colour concepts turn out to be invariant through various cultures, and they correspond rather precisely to the different types of colour sensors (the cones) in the human retina. In the domain of taste, there is a good correspondence between basic level concepts (sweet, salty, sour, bitter…) and the structure of taste sensors in the human tongue. On the basis of data on the visual cortex of monkeys, we can expect that part of categorisation abilities concerning hands (or faces) is hardwired in our brain, and so on.

Something similar also happens in the case of categories that do not have a specific perceptual nature, as in the case of emotions, where basic level categories roughly correspond to so-called basic emotions (anger, joy, fear, surprise...), which are common to humans and other mammals, and which are therefore presumably innate and depend on deep features of the cognitive architecture.

Therefore, the divide is not between "abstract" and perceptual concepts (or concepts depending on a specific sensorial modality). Another example of an abstract (and not modality-specific) concept that seems to be (at least

partially) localised in some specific cognitive structure relates to (natural) numbers.

If the advocates of massive modularity were right, there would be many examples of concepts that do not pertain to some specific perceptual modality, and are at the same time localised in some specific structure of the cognitive architecture, Examples would be concepts pertaining to naïve physics, naïve biology, naïve psychology, and so on. However, it is not necessary to accept the thesis of massive modularity to conclude that (even rather abstract) concepts may be grounded in very different cognitive subsystems. It is sufficient to accept the (much less controversial) assumption that within the mind there are some "dedicate" components that are responsible for some of our categorisation abilities.

The stance of massive modularism, and the thesis that some conceptual abilities depend on specific cognitive mechanisms has seldom been associated with the thesis of the domain specificity of concepts (see e.g. Pacherie 1993). Note, however, that even within the same domain, concepts may be represented in deeply different ways. Consider the domain of natural numbers. According to various kinds of empirical evidence, "small" numbers and "large" numbers are mentally represented in rather different ways. According to (Dehaene 1997) for example, small numbers correspond to an innate, analogue representation (the so called *accumulator*), which is common to humans and other animal species. The representation of large numbers instead depends on rather distributed, high level abilities. Therefore, from the point of view of corresponding mental representations, the apparent uniformity of number concepts is misleading.

Recently Lawrence Barsalou (2005) advocated a view of the conceptual system as "non-modular and modal", in that conceptual representations would be scattered over different modality-specific areas. Starting from Gil-da-Costa experimental results, he takes into account the knowledge of scream situations in macaques, which involves different modality-specific systems (auditory, visual, affective systems...). Barsalou interprets this data in favour of the thesis of a continuity of conceptual representations in different animal species, in particular between humans and non-human mammals: "this same basic architecture for representing knowledge is present in humans. [...] knowledge about a particular category is distributed across the modality-specific systems that process its properties" (p. 309). I find such considerations plausible and convincing. However, I wonder in which sense we can still speak of a conceptual *system* here. If it is intended in this way, the conceptual system coincides with the whole mind.

Summing up, at one extreme there are “concepts” that have a precise counterpart in localised aspects of the cognitive system – however, such aspects are scattered over the entire mental architecture. Other “concepts” are more holistic in nature, being dependent in a crucial way on inferential networks, and are built on learning, experience, culture, and so on. Some “concepts” are presumably “made of words”, and exist only within some linguistic game.

In conclusion, it is likely that a collection of homogeneous mental structures corresponding to the ordinary notion of concept does not exist. In other words, it is reasonable to doubt that concepts form a “natural kind” from the viewpoint of cognitive science. A similar thesis has been proposed by Edouard Machery (2005), who argues for the existence of three kinds of theoretical constructs, corresponding to the three approaches mentioned in section 1 (prototype view, exemplar view, theory theory). In my opinion, however, more radical conclusions can be drawn: Categorisation tasks are carried out in different ways by different components of the cognitive architecture. In some cases, the mind categorises the world in a “holistic” way; in other cases, categorisation is performed by some localised component in the mental architecture. In all such cases, some regularity emerges from categorisation behaviour (first of all, prototypical effects). But this is due to the fact that the mind must face similar constrains, rather than to some common underlying cognitive mechanism.

References

Armstrong, S.L., Gleitman, L.R. & Gleitman, H. (1983). What some concepts might not be, *Cognition*, 13: 293-308.

Barsalou, L.W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure of categories, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11: 629–654.

—. (1987). The instability of graded structures: Implications for the nature of concepts, in U. Neisser (ed.), *Concepts and Conceptual Development*, Cambridge University Press, Cambridge, UK.

—. (2005). Continuity of the conceptual system across species, *Trends in Cognitive Science*, 9(7): 305-311.

Connolly, A.C., Fodor, J.A., Gleitman, L.R. & Gleitman, H. (2007). Why stereotypes don’t even make good defaults, *Cognition*, 103, 1-22.

Dehaene, S. (1997). *The Number Sense*, Oxford University Press.

Fodor, J. & Lepore, E. (1996). The red herring and the pet fish: why concepts still can't be prototypes, *Cognition*, 58: 253-270.

Hampton, J.A. (1999). Concepts, in R.A. Wilson & F.C. Keil (eds.), *The MIT Encyclopedia of the Cognitive Sciences*, The MIT Press, Cambridge, MA.

Harris, M. (1985). *Good to Eat. Riddles of Food and Culture*, Simon and Schuster, New York

Koons, R.C. (2005). Defeasible reasoning, *Stanford Encyclopedia of Philosophy* (http://plato.stanford.edu/).

Laurence, S. and Margolis, E. (1999). Concepts and Cognitive Science, in Margolis and Laurence (1999).

Machery, E. (2005). Concepts are not a natural kind, *Philosophy of Science*, 72 (2005), 444–467.

Margolis, E. & Laurence, S. (1999). *Concepts. Core Readings*, The MIT Press, Cambridge, MA.

Medin, D.L., Wattenmaker, W.D. & Hampson, S.E. (1987). Family resemblance, conceptual cohesiveness, and category construction, *Cognitive Psychology*, 19, 242-279.

Miller, G.A. & Johnson-Laird, Ph.N. (1976). *Language and Perception*, Harvard University Press, Cambridge, MA.

Minsky, M. (1975). A framework for representing knowledge, in P. Winston (ed.), *The Psychology of Computer Visio*n, McGraw-Hill, New York.

Murphy, G.L. (2002). *The Big Book of Concepts*, The MIT Press, Cambridge, MA.

Pacherie, E. (1993). L'hypothèse de la structuration des connaissances par domaines et la question de l'architecture fonctionnelle de l'esprit, *Revue Internationale de Psychopathologie*, 9, pp. 63-89.

Putnam, H. (1975). *Mind, Language and Reality*, Cambridge University Press, Cambridge.

Rey, G. (1983). Concepts and stereotypes, *Cognition*, 15: 237-262.