Concepts, perception and the dual process theories of mind

Marcello Frixione[1], Antonio Lieto[2]

[1] DAFIST, University of Genova, via Balbi, 4 - 16126 Genova, Italy

[2] Dept. of Computer Science, University of Torino, C.So Svizzera, 185 - 10149 Torino, Italy

and

ICAR - CNR, Viale delle Scienze Ed. 11, 90128 Palermo, Italy

[1]*marcello.frixione@unige.it*, [2]*lieto@di.unito.it*

Abstract

In this article we argue that the problem of the relationships between concepts and perception in cognitive science is blurred by the fact that the very notion of concept is rather confused. Since it is not always clear exactly what concepts are, it is not easy to say, for example, whether and in what measure concept possession involves entertaining and manipulating perceptual representations, whether concepts are entirely different from perceptual representations, and so on. As a paradigmatic example of this state of affairs, we will start by taking into consideration the distinction between conceptual and nonconceptual content. The analysis of such a distinction will lead us to the conclusion that *concept* is a heterogeneous notion. Then we shall take into account the so called dual process theories of mind; this approach also points to concepts being a heterogeneous phenomenon: different aspects of the conceptual competence are likely to be ascribed to different types of systems. We conclude that without a clear specification of what concepts are, the problem of the relationships between concepts and perception is somewhat ill-posed.

1. Introduction

In our opinion, the problem of the relationships between concepts and perception in cognitive science is blurred by the fact that the very notion of concept is rather confused. Since it is not always clear exactly what concepts are, it is not easy to say, for example, whether and in what measure concept possession involves entertaining and manipulating perceptual representations, whether concepts are entirely different from perceptual representations, and so on. As a paradigmatic example of this state of affairs, we will start with some considerations on the distinction between conceptual and nonconceptual content (sect. 2). The analysis of such a distinction will lead us to the conclusion that *concept* is a heterogeneous notion. In sect. 3 we shall take into account the so called dual process theories of mind. This approach also indicates that concepts are likely to be a heterogeneous phenomenon: it is plausible that different aspects of conceptual competence must be ascribed to different types of systems. Section 4 illustrates an example in this sense: compositionality on the one hand and (some aspects of) typicality effects on the other can be accounted for in terms of different types of representation. Some conclusions follow (sect. 5).

2. Conceptual vs. Nonconceptual Content

The notion of nonconceptual content was initially proposed by Gareth Evans (1982), and subsequently adopted by many other philosophers, such as José Luis Bermudez, Tim Crane and Christopher Peacocke. Bermudez & Cahen (2011) is an overview of the debate on this topic, and Gunther (2003) is a collection of classical papers on the issue. The central idea is that some mental states have a representational content that is not structured in terms of concepts. In this section we argue that the viability of the notion of nonconceptual content (and, in particular, its profitability for cognitive science) is undermined by the fact that: i) the use of the term "concept" in philosophy is often not homogeneous with its use in cognitive psychology, and ii) even within the fields of philosophy and psychology considered separately, a coherent notion of concept does not emerge.

According to nonconceptualist philosophers, nonconceptual contents do not require conceptual abilities. As a consequence, if some mental state is nonconceptual, its bearer may not even possess the concepts needed to specify its content. Initially, this notion was applied to the phenomenal content of perceptual states; subsequently it was extended to other fields (such as subpersonal and computational states). Philosophers are currently debating whether or not such nonconceptual contents exist, or if mental content is always organized in terms of concepts (as is claimed by McDowell (1994), for example). Within cognitive science, proposals for applying the conceptual/nonconceptual distinction have been developed for example by Raftopoulos and Müller (2006), and, in the field of visual perception, by Jacob and Jeannerod (2003). During the 1990s the notion of nonconceptual content was applied to representations in connectionist models (Cussins 1990; Clark 1994).

In our opinion some aspects of the conceptual/nonconceptual distinction appear problematic when applied to cognitive science (Dell'Anna e Frixione 2010). As discussed in recent philosophical literature (for instance see Machery 2005, 2009; Piccinini 2011), difficulties arise in characterizing the notion of concept itself. In the first place, the use of the term "concept" in the philosophical tradition is often not homogeneous with the use of the same term in empirical psychology. Briefly, we could say that in cognitive psychology the emphasis is on such tasks as categorization, learning or induction, and a concept is essentially intended as the mental representation of a category. According to many philosophers, on the other hand, concepts are above all the components of thoughts. Even if we leave aside the problem of specifying exactly what thoughts are, this requires a more demanding notion of concept. In other words, some phenomena that are classified as conceptual by psychologists turn out to be nonconceptual for philosophers. There are, thus, mental representations of categories that philosophers would not consider genuine concepts. For example, according to many philosophers, concept possession involves the ability to make explicit, high level inferences, and sometimes also the ability to justify them (Peacocke 1992; Brandom 1994; Bermudez 1995). This clearly exceeds the possession of the mere mental representation of categories. Moreover, according to some philosophers, concepts can be attributed only to agents who can use natural language (i.e. only adult human beings).

On the other hand, psychologists are more tolerant on concept attribution. Elizabeth Spelke's experiments on infants (see e.g. Spelke 1994; Spelke and Kinzler

2007) are symptomatic of the difference in approach between (some) psychologists and (some) philosophers. Such experiments demonstrate that the mental representation of some extremely general categories is very precocious and presumably innate. According to the author, these experiments show that very young children possess certain *concepts* (e.g. the concept of physical object). But among nonconceptualist philosophers, such as Bermudez (Bermudez 1995, Bermudez & Cahen 2011), these same data have been interpreted as a paradigmatic example of the existence of nonconceptual contents in agents (children) that have not yet developed a conceptual system.

The fact that philosophers often consider concepts mainly as the components of thoughts has led to a great emphasis on compositionality and on related features, such as productivity and systematicity, that are often less important in the psychological study of concepts. On the other hand, it is well known that compositionality is at odds with typicality effects, which are crucial in most psychological characterizations of concepts (see sect. 4 below). The distinction between conceptual and nonconceptual content has been developed within the philosophical tradition, and therefore rests on an essentially "philosophical" notion of concept. Given the state of affairs described above, one could suspect that such a distinction cannot be profitably integrated into the research on concepts in cognitive psychology and, more generally, into the discourse of cognitive science. This, of course, should not be a problem if the philosophical notion of concept were clear and theoretically useful. Nevertheless, there remains the problem of conciliating the philosophical and the psychological notions of concept in order to avoid conflicts and confusions.

Unfortunately, things are made more complex by the fact that, even within the two fields of research considered separately, the situation is not much more encouraging. In neither of the two disciplines does a clear, unambiguous and coherent notion of concept seem to emerge. In psychology there are different positions and theories on the nature of concepts (prototype view, exemplar view, theory theory, etc.) that cannot be easily be integrated. From this point of view, the conclusions of Murphy (2002) are of great significance, since in many respects this book reflected the status of empirical research on concepts. Murphy contrasts the approaches mentioned above in relation to different classes of problems, including learning, induction, lexical concepts and children's concepts. His conclusions are worrying: the result of comparing the various approaches is that ''there is no clear, dominant winner'' (ibid. p. 488) and that ''[i]n short, concepts are a mess'' (p. 492). This situation persuaded some scholars to doubt whether concepts constitute a homogeneous phenomenon from the point of view of a science of the mind (see e.g. Machery 2005, 2009). The situation is also rather discouraging among philosophers. The disagreement even concerns the ontological status of concepts. Some philosophers accept a mentalistic position, according to which concepts must be identified with some class of mental phenomena. Others share a more traditional anti-psychologistic stance, and refuse to identify concepts with mental states or processes. For example (considering authors who accept the existence of nonconceptual content) Bermudez & Cahen (2011) states that ''concepts are semantic entities rather than psychological entities'', and on this point they agree with Peacocke (1992). Conversely, Michael Tye (2006) claims that ''[concepts] are mental representations of a sort that can occur in thought'' (p. 506). The mental nature of concepts is accepted by philosophers like Fodor and Dretske among others. A disagreement of such significance casts doubt on whether we are in fact dealing with

different notions. And, above all, one wonders what utility there can be for the empirical psychology of anti-psychologistic positions that deny the mental nature of concepts. However, disagreement is not limited to this point. For example, those who claim that concepts must be characterized in terms of (some or all of) the inferences in which they are involved are inclined to accept some form of holistic or molecularist position that is strongly opposed by supporters of atomism, first and foremost Jerry Fodor (1998). Furthermore, another problem that emerges when considering the conceptual/non conceptual distinction is that the notion of non conceptual content likewise seems to be very heterogeneous. For example Tye claims that the nonconceptual content of visual perception encompasses everything that falls within the scope of foveal vision. Yet, at the same time, the phenomena studied by Elizabeth Spelke in her experiments should also concern nonconceptual content. In other words, nonconceptual content should include both raw proximal stimuli and representations in which data are processed and undergo complex forms of categorization (for example in terms of physical objects, faces or cause-effect relations). Consider also Jacob and Jeannerod's (2003) model, according to which visual percepts and visuomotor representations (i.e. the representations processed respectively by the ventral and the dorsal path of the visual system) are both endowed with nonconceptual content, even though these two kinds of representation are profoundly different from the point of view both of their format and their purpose.

3. Concepts and Dual Process Theories

Summing up, there are good reasons to suspect that the very notion of concept is somewhat heterogeneous, and that this negatively affects the possibility of conclusively investigating such distinctions as conceptual vs. nonconceptual, conceptual vs. perceptual, and so on. We have already mentioned that in psychology different positions and theories on the nature of concepts have been proposed; these are usually grouped into three main classes: prototype views, exemplar views and theory-theories (see again Murphy 2002). These approaches turn out not to be mutually exclusive. Rather, they seem to succeed in explaining different classes of cognitive phenomena, and many researchers hold that all of them are needed to explain psychological data. This consideration led Edouard Machery (2005, 2009) to claim that *concept* is not a natural kind, hypothesising that three different natural kinds exist, corresponding respectively to prototypes, exemplars and theories. If concepts are not a homogeneous category, then the problem of the relationships between concepts and perception must be split into different problems. In the following we advance the hypothesis that there are other reasons to suspect that concepts are not a homogeneous class of entities from the standpoint of cognitive science. In particular, we shall take into account suggestions from the so called dual process theories of mind. As we shall see, this should also influence the problem of the relationships between concepts and perception.

According to the *dual process theories* (Stanovich & West 2000, Evans & Frankish 2008, Kahneman 2011) two different types of cognitive processes and systems exist, which have been called respectively system 1 and system 2.

System 1 processes are automatic. They are phylogenetically older, and are shared by humans and other animal species. They are innate, and control instinctive behaviors, so

4

they do not depend on training or specific individual abilities and, in general, are cognitively undemanding. They are associative, and operate in a parallel and fast way. Moreover, system 1 processes are not consciously accessible to the subject.

System 2 processes are phylogenetically more recent, and are specific to the human species. They are conscious and cognitively penetrable (i.e. accessible to consciousness), and are based on explicit rule following. As a consequence, if compared to system 1, system 2 processes are sequential and slower, and cognitively more demanding. Performances that depend on system 2 processes are usually affected by acquired skills and differences in individual capabilities.

The dual process approach was originally proposed to account for systematic errors in reasoning: systematic reasoning errors (consider the classical examples of the selection task or the conjunction fallacy) should be ascribed to fast, associative and automatic system 1 processes, while system 2 is responsible for the slow and cognitively demanding activity of producing answers that are correct with respect to the canons of normative rationality.

In our opinion, the distinction between system 1 and system 2 processes may also be plausibly applied to the problem of conceptual representations as emerging in the above section (for a similar position in this respect, see Piccinini 2011; we shall briefly return to Piccinini's stance in the conclusions of this paper – sect. 5). For example, categorization based on typical traits (either prototype based or exemplar based) is, presumably, in many cases a fast and automatic process which does not require any explicit effort, and which therefore could presumably be attributed to a type 1 system. On the contrary, there are types of inference that are usually included within "conceptual competence", which are slow and cognitively demanding and which should be attributed to processes that are more likely to be ascribed to type 2. An example could be categorization processes based on complex classical definitions given in terms of necessary and sufficient conditions (indeed, according to the so-called classical theory, concepts can be defined in terms of sets of individually necessary and jointly sufficient conditions).

Consider for example the well-known Linda problem and the already mentioned conjunction fallacy (Tversky & Kahneman 1983). In a well known experiment from the psychology of probabilistic reasoning, subjects are given a description of a person named Linda that perfectly fits the stereotype of a feminist activist. Then they are asked to judge if it is more likely that Linda is (a) a bank teller or (b) a bank teller *and* a feminist. The majority of subjects choose (b), without realizing that being a feminist bank teller is in any case more demanding (and therefore less probable) than simply being a bank teller. Indeed, the conjunction fallacy consists in failing to realize that a conjunction is always less probable than its conjuncts. From the standpoint of a theory of reasoning, this is a paradigmatic case of a systematic error that, according to the dual process approach, can be explained in terms of the system 1 vs. system 2 distinction: the conjunction fallacy is an exemplary case of the effects of a system 1 process. However, if we consider the problem from the point of view of a theory of concepts, the conjunction fallacy can be interpreted as an example of the strong tendency of human subjects to resort to prototypical information in categorization, even when this is not appropriate (at least from the point of view of those who devised the experiment): Linda

is categorized as a feminist because she perfectly fits the prototypical traits of a feminist (while she does not in the least fit the prototypical traits of a bank teller). So, we could see the "error" as being originated by a process of categorization based on prototypical knowledge.

Typicality effects in categorisation and, in general, in category representation are crucial for human cognition. Under what conditions should we say that somebody *knows* the concept DOG (or, in other terms, that (s)he possesses an adequate mental representation of it)? It is not easy to say. However, if a person does not know that, for example, dogs usually bark, that they typically have four legs and that their body is covered with fur, that in most cases they have a tail and that they wag it when they are happy, then we probably should conclude that this person does not grasp the concept DOG. Nevertheless, all these pieces of information are neither necessary nor sufficient conditions for being a dog. In fact, they are traits that characterise dogs in typical (or prototypical) cases.

The concept DOG is not exceptional from this point of view. The majority of everyday concepts behave in this way. For most concepts, a classical definition in terms of necessary and sufficient conditions is not available (or, even if it is available, it is unknown to the agent). On the other hand, it may be that we know the classical definition of a concept, but typical/prototypical knowledge still plays a central role in many cognitive tasks. Consider the following example: nowadays most people know the necessary and sufficient conditions for being WATER: water is the chemical substance whose formula is H_2O, i.e., the substance whose molecules are formed by one atom of oxygen and two atoms of hydrogen. However, in most cases of everyday life, when we categorise a sample of something as WATER, we do not take advantage of this piece of knowledge. We use such typical traits such as the fact that (liquid) water is usually a colourless, odourless and tasteless fluid.

The use of typical knowledge in cognitive tasks such as categorisation has to do with the constraints that concern every finite agent that has a limited access to the knowledge relevant for a given task. In most cases, cognitive processes based on typical knowledge are fast and automatic, cognitively undemanding, and are presumably homogenous to the processes employed in similar tasks by non-human animals. Consider for example the following variant of the Linda problem. Let us suppose that a certain individual Pippo is described as follows. He weighs about 200 kg, and he is approximately two meters tall. His body is covered with a thick, dark fur, he has a large mouth with robust teeth and paws with long claws. He roars and growls. Now, given this information, we have to evaluate the plausibility of the two following alternatives:

> a) Pippo is a mammal;
> b) Pippo is a mammal, and he is wild and dangerous.

Which is the "correct" answer? According to the dictates of the normative theory of probability, it is surely a). But if you encounter Pippo in the wilderness, it would probably be best to run.

So, many aspects of the psychology of concepts have presumably to do with fast, type 1 system and processes, while others can be more plausibly ascribed to type 2. In particular, the ability to make explicit, high level inferences, and, moreover, the ability to justify them, which some philosophers consider to be constitutive of concept

possession (see, for instance, Peacocke 1992; Brandom 1994), can be plausibly considered as type 2 processes.

Problems arise from the fact that the dual process approach is not monolithic. Different dual process theories exist – a detailed review is provided in Evans (2008). However, the dual process approach has also received a number of criticisms, many of which have recently been reviewed and answered by, Evans and Stanovich (2013a, b). (For a skeptical attitude towards the possibility of applying the dual process theories to the study of concepts see sect. 8 of Machery 2011.) In any case, it is likely that many kinds of type 1 systems and processes exist, with partially different properties. Consider, for example, expertise: an art historian can distinguish a painting by, say, Rubens from one by van Dyck at a glance and without the need for any form of conscious, sequential application of explicit rules. But surely this ability is not innate and depends on specific training (and, presumably, it is not shared with other animal species). A recurring distinction in the dual process literature concerns the interaction of the two types of system. A first possibility (Sloman 1996) is that they operate in parallel. In this case, adopting the terminology proposed by Evans (Evans 2008), the two systems are assumed to be parallel-competitive. A second possibility (Kahneman & Frederick 2002, Kahnemann 2011, Evans and Stanovich 2013a) is the so called default-interventionist approach. According to this view, deliberative type 2 reasoning processes can inhibit the possibly biased, default responses of type 1 systems, and replace them with the outputs of reflective reasoning. This second perspective better fits our approach.

However, our concern here is not to take a stand for some specific version of the theory. In particular, we do not make any claim about how many systems involved in the two types of process effectively exist. Rather, our claim is that dual process theories can supply some useful suggestions in order to understand and classify the wide class of cognitive phenomena that pertain to human conceptual abilities.

A clarification is appropriate. The appeal to dual process theories could be interpreted as a way to re-introduce some form of the conceptual/nonconceptual distinction. This is definitely not our intent. In spite of some superficial similarities, the system 1 vs. system 2 distinction does not overlap with the conceptual vs. non conceptual opposition. In the first place, the conceptual/nonconceptual dichotomy originated in a philosophical context: it rests on conceptual, a priori analyses and was proposed while disregarding any form of evidence coming from the empirical study of the mind. Only later did some theorists try to apply it to cognitive science, but this, far from clarifying the situation, generated further confusion (for example, it is not even clear if nonconceptual content must be situated at a personal or a subpersonal level of analysis). Moreover, the conceptual vs. nonconceptual distinction is between two types of *content*, and the philosophical notion of content is a semantic notion that is not immune from problems, and it is not clear if and in what measure it can be profitably adopted within the scope of an empirical science of mind. On the other hand, the system1 vs. system2 dichotomy is an empirical distinction that originated within the field of cognitive science, and it does not concern different types of content, but different types of processes and/or systems that can be individuated within the mental architecture. Finally, the conceptual/nonconceptual distinction presupposes a notion of concept, and, in our opinion, the point is precisely that a clear and unproblematic notion of concept is lacking. The system 1/system 2 opposition has the advantage of having

been developed independently; therefore, the debate concerning dual process approaches should offer a neutral point of view in order to evaluate some problems concerning concepts in cognitive science.

Compositionality, which is often considered to be an irrevocable trait of conceptual systems, could turn out to be more akin to system 2 abilities. Compositionality has to do with higher cognition and with complex inferential tasks: paradigmatic examples of compositional systems are natural languages and logical formalisms. And compositionality is somewhat at odds with typicality effects: the characterisation of concepts in typical terms is difficult to reconcile with the requirement of compositionality. According to a well-known argument (Fodor 1981; Osherson & Smith 1981) prototypes are not compositional. In brief, the argument runs as follows: consider a concept like PET FISH. It results from the composition of the concept PET and of the concept FISH. However, the prototype of PET FISH cannot result from the composition of the prototypes of PET and FISH: a typical PET is furry and warm, a typical FISH is greyish, but a typical PET FISH is neither furry and warm nor greyish. A possible solution should be to hypothesize that compositional representations and representations in (proto)typical terms depend on different cognitive components, based on different types of representation.

In this respect it may be interesting to take into account what happens in the field of the computational simulations of cognition (Frixione & Lieto 2012). In artificial intelligence, the representation of prototypical information is problematic in compositional knowledge representation formalisms. For example, *description logics* (Baader et al. 2010) are a widespread class of concept oriented representation systems (the Web Ontology Language OWL belongs to this class). They allow the representation of taxonomies of concepts in terms of sets of necessary and/or sufficient conditions, but do not allow for the representation of typical traits. On the other hand, description logics are fully compositional systems: they are subsets of first order predicate logic, and logical formalisms are compositional. Consider *classification*, one of the most characteristic forms of inference defined on this type of formalism: classifying a concept amounts to individuating its more specific superconcepts and its more general subconcepts, or, in other words, to identify implicit superconcept-subconcept relations in a taxonomy. For human subjects such a process is far from natural, fast and automatic: it is usually slow, it can require great effort and it is facilitated by specific training. So, in terms of dual process theories, the inferential task of classifying concepts in taxonomies is *prima facie* a type 2 process, qualitatively different from the task of categorizing items as instances of a certain class on the basis of typical traits (e.g. the task of categorizing Fido as a dog because he barks, has fur and wags his tail). Therefore, it is plausible that conceptual representation in computational systems should be assigned to (at least) two different kinds of components responsible for different tasks: type 2 processes, involved in complex and cognitively demanding inference tasks, and fast and automatic type 1 processes (such as those involved in categorization based on prototypical information). In the next section we shall explore some aspects of this hypothesis.

4. Compositionality, Typicality and Conceptual Spaces

Let us consider an argument against the possibility of reconciling compositionality and typicality effects that dates back to Osherson and Smith (1981), and which is a version of the pet fish argument mentioned above. Osherson and Smith's original aim was to show that fuzzy logic is inadequate to capture typicality. However, as we shall see, their argument has a wider range of application.

At first sight, fuzzy logic seems to be a promising approach in order to face the problem of typicality. Indeed, one consequence of typicality effects is that some members of a category C turn out to be better (i.e. more typical) instances of C than others. For example, a robin is a better example of the category of birds than, say, a penguin or an ostrich. More typical instances of a category are those that share a greater number of typical features (e.g. the ability to fly for birds, having fur for mammals, and so on). The fuzzy value of a predicate (say, F) could be interpreted as a measure of typicality: given two individuals h and k, it seems natural to assume that $F(h) > F(k)$ iff h is a more typical instance of F than k is.

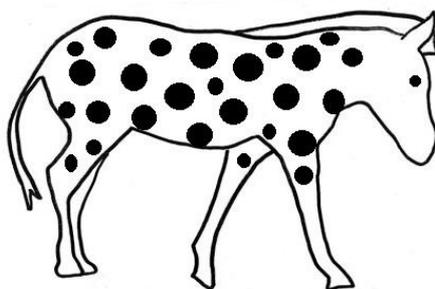However, let us consider the zebra in fig. 1 (and let us suppose that her name is Pina).

Fig. 1

Pina is presumably a good instance of the concept POLKA DOT ZEBRA[1]; therefore, if such a concept were represented as a fuzzy predicate, then the value of the formula *polka_dot_zebra*(*Pina*) should be close to 1, say:

(1) *polka_dot_zebra*(*Pina*) = .97

On the other hand, Pina is a rather poor (i.e. atypical) instance of the concept ZEBRA; therefore the value of the formula *zebra*(*Pina*) should be low, say:

(2) *zebra*(*Pina*) = .2

(of course, the specific values are not relevant here; the point is that Pina is more typical as a polka dot zebra than as a zebra). But POLKA DOT ZEBRA can be expressed as the conjunction of the concepts ZEBRA and POLKA DOT THING; i.e. in logical terms, it holds that:

(3) $\forall x$ (*polka_dot_zebra*(*x*) $\leftrightarrow$ *zebra*(*x*) $\wedge$ *polka_dot_thing*(*x*))

---

[1] Osherson and Smith's original example was not a polka dot zebra but a striped apple.

Now, the problem is the following: if we adopt the simplest and most widespread form of fuzzy logic, the value of a conjunction is calculated as the minimum of the values of its conjuncts, and this makes it impossible for the value of *zebra*(*Pina*) to be .2 and that of *polka_dot_zebra*(*Pina*) to be .97 at the same time. Of course, there are other types of fuzzy logic, in which the value of a conjunction is not the minimum of the values of the conjuncts. But in no case can a conjunction exceed the value of its conjuncts. Worse still, in logic, once a suitable order has been imposed on truth values, it generally holds that:

$$val(A \wedge B) \leq val(A) \text{ and } val(A \wedge B) \leq val(B)$$

So, the problem pointed out by Osherson and Smith does not seem to concern only fuzzy logic. Rather, Osherson and Smith's argument seems to show that, in general, logic based representations are unlikely to be compatible with typicality effects. And, as mentioned before, logic based representations are paradigmatic examples of compositional systems, which fully embody the Fregean principle of compositionality of meaning. (Note also that this is exactly the same problem which, in the context of probabilistic reasoning, gives rise to the conjunction fallacy mentioned in sect. 3 above.)

Indeed, the situation seems to be more promising if, instead of logic, we face typicality by adopting some other form of representation, such as for example conceptual spaces, a geometric representation proposed by Peter Gärdenfors (2000). A *conceptual space* (CS) is a space in a certain number of quality dimensions. Concepts are represented in the terms of such dimensions, and correspond to regions in CSs. CSs dimensions can be more or less directly related to perception (such as temperature, weight, brightness, pitch), or more abstract in nature. Each quality dimension is associated with a geometrical (topological or metrical) structure. The central idea behind this approach is that knowledge representation takes advantage of the geometrical structure of conceptual spaces. For example, similarity should be calculated in terms of some suitable distance measure. So, if we represent a concept as a (convex) area in a suitable conceptual space CS, then the degree of typicality of a certain individual can be measured as the distance of the corresponding point from the centre of the area. The conjunction of two concepts is represented as the intersection of the two corresponding areas, as in fig. 2.
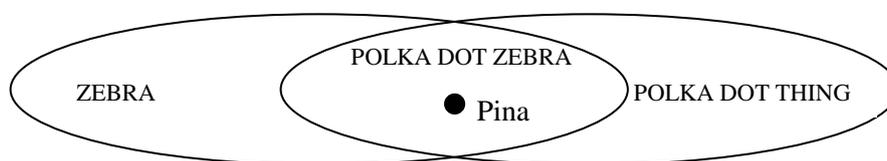

Fig. 2

According to the conceptual space approach, Pina should presumably turn out to be very close to the centre of POLKA DOT ZEBRA (i.e. to the intersection between ZEBRA and POLKA DOT THING). In other words, she should turn out to be a very typical polka dot zebra, despite being very eccentric with respect to both the concepts ZEBRA and

POLKA DOT THING (that is to say, she is an atypical zebra and an atypical polka dot object). This seems to better capture our intuitions about typicality. We should conclude that the treatment of compositionality and that of (some forms of) typicality require rather different approaches and forms of representation, and should therefore presumably be assigned to different components of the cognitive architecture.

The above considerations can be reconciled with both a hybrid approach to concepts (Miller & Johnson-Laird 1976, Osherson & Smith 1981) and with the so-called heterogeneity hypothesis (Machery & Seppälä 2010, Machery 2014). Here we do not take a stand on this point. Our concern is simply to stress that some aspects of typicality effects have features that are closer to those usually associated to type 1 systems and processes (they are fast, automatic, and so on), while compositional representations seem to better fit type 2 tasks. This favours the hypothesis of adopting different formalisms for representing different aspects of conceptual knowledge.

5. Some Conclusions

In conclusion, different aspects of "conceptual competence" seem to involve different types of cognitive processes (e.g. type 1 vs. type 2 processes), and seem to require different kinds of representation (e.g. compositional, "linguistic" representations vs. some other type of representation, as conceptual spaces). Perception (or, at least, *low-level* perception) is presumably more akin to type 1 processes, and more remote from type 2 ones. Therefore, those aspects of concepts (if any) that are related to perception presumably pertain to type 1 processes. Moreover, such geometric representations as conceptual spaces seem prima facie to be closer to the format of the output of perceptual systems.

We do not maintain that typicality pertains exclusively to fast, low level type 1 processes. Certain forms of categorization, which are certainly more akin to type 2 process, do not rest on classical definitions given in terms of necessary and sufficient conditions. Let us consider the following hypothetical example: suppose that a rather inexperienced amateur naturalist (let us call him Jean-Baptiste) finds a specimen like that in fig. 3 A[2].
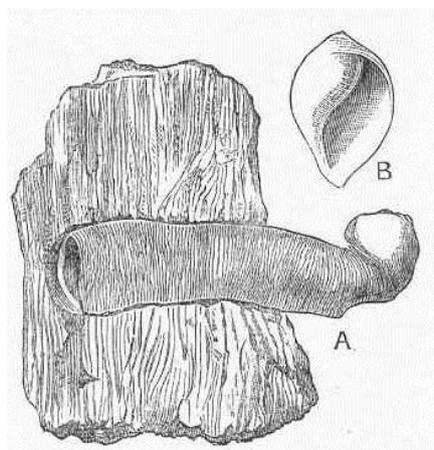

Fig. 3

Prima facie, Jean Baptiste classifies this calcareous twisted tube, embedded in a piece of coral, as the shell of a worm of the phylum Annelida. However, upon closer analysis, he notes that the shape of the initial portion of the tube is similar to the spire of a gastropod (a "snail"). Now, Jean-Baptiste knows various things about gastropoda; for example, he knows that the apical spires of gastropod shells are the first to develop. So, when the animal in fig. 3 A was young, it presumably resembled a small snail (as in fig. 3 B), which later clung to a solid substrate and developed "as a worm". Jean-Baptiste knows also that the juvenile shape of a life form is very important in order to classify it. So, he hypothesizes that the twisted tube in fig. 3 A is actually a mollusk shell, albeit a strange one. (And he is right, because it effectively is a gastropod of the genre *Magilus*.)

Now, a process like this certainly cannot be ascribed to a type 1 system: it is slow, sequential, penetrable to consciousness, and it depends on high level, acquired pieces of knowledge. But neither does it resemble the application of a classical definition; it rather depends on forms of typicality-based reasoning, on abductions and on "theory based" inferences.

As said before, Machery (2005, 2009) claims that *concept* is not a natural kind; he rather hypothesises three different kinds, which correspond respectively to prototypes, exemplars and theories. Piccinini (2011) argues for an alternative point of view, according to which only two kinds of concept exist: implicit and explicit; he correlates implicit and explicit concepts respectively to system 1 and system 2 processes. In our opinion, it is likely that in some respect both Machery and Piccinini are right, in that they both individuate important discontinuities in the traditional notion of concept. However, it is also likely that the distinction between system 1 vs. system 2 processes is in part orthogonal to Machery's tripartition, and there good reasons to suspect that they are not mutually exclusive (Frixione and Lieto 2012, Frixione 2013).

In conclusion, the situation is rather complex. It is plausible that, in light of the distinction between type 1 vs. type 2 systems, certain aspects of conceptual knowledge pertain to type 1 systems, and that "type 1 concepts" are likely to be closer to perceptual processes. But, without a clear specification of what concepts are, the problem of the relationships between concepts and perception remains somewhat ill-posed; a satisfactory solution would presuppose a better understanding of the notion of *concept* itself.

References

Baader, F., D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (2010). *The Description Logic Handbook: Theory, Implementations and Applications*, 2nd edition. Cambridge, UK: Cambridge University Press.

Barsalou, L.W. (1985). Continuity of the conceptual system across species. *Trends in Cognitive Science*, 9(7), 305-311.

Bermudez, J.L. (1995). Nonconceptual content: from perceptual experience to subpersonal computational states. *Mind and Language*, 10, 333-369. Reprinted in Gunther (2003).

Bermudez, J.L. & Cahen, A. (2011). Nonconceptual mental content. *Stanford Encyclopaedia of Philosophy*, http://plato.stanford.edu/

Brandom, R. (1994). *Making it Explicit*. Cambridge, MA: Harvard University Press.

Clark, A. (1994). Connectionism and cognitive flexibility. In T. Dartnall (Ed.), *Artificial Intelligence and Creativity* (pp. 63-79), Dorderecht: Kluwer Academic Publishers. Reprinted in Gunther (2003).

Cussins, A. (1990). The connectionist construction of concepts. In M. Boden (Ed.), *The Philosophy of Artificial Intelligence* (pp. 380-400). Oxford: Oxford University Press. Partially reprinted in Gunther (2003).

Cussins, A. (2003). Postscript to the partial reprint of Cussins (1990). In Gunther (2003), 147-159.

Dell'Anna, A., Frixione, M., (2010). On the advantage (if any) and disadvantage of the conceptual/nonconceptual distinction for cognitive science. *Minds & Machines*, 20, 29-45.

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

Evans, J.St.B.T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.

Evans, J.St.B.T. & Frankish, K.E. (eds.) (2008). *In Two Minds: Dual Processes and Beyond*. New York, NY: Oxford UP.

Evans, J.St.B.T. & Stanovich, K.E. (2013a). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science*, 8, 223-241.

Evans, J.St.B.T. & Stanovich, K.E. (2013b). Theory and Metatheory in the Study of Dual Processing: Reply to Comments, in *Perspectives on Psychological Science*, 8, 263-271.

Fodor, J. (1981). The present status of the innateness controversy. In J. Fodor, *Representations*, Cambridge, MA: MIT Press.

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

Frixione, M. (2013). Concepts and Fat Plants: Non-Classical Categories, Typicality Effects, Ecological Constraints. *Concepts – Contemporary and Historical Perspectives*, *ProtoSociology*, 30, 152-166.

Frixione, M. & Lieto, A. (2012). Representing Concepts in Formal Ontologies. Compositionality vs. Typicality Effects. *Logic and Logical Philosophy*, 21(4), pp. 391–414. DOI: 10.12775/LLP2012.018.

Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*, MIT Press, Bradford Books.

Gunther, Y.H. (Ed.) (2003). *Essays on Nonconceptual Content*. Cambridge, MA: The MIT Press.

Jacob, P. & Jeannerod, M. (2003). *Ways of Seeing. The Scopes and Limits of Visual Cognition*. Oxford, UK: Oxford University Press.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.

Kahneman, D. & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment*, pp. 49–81. Cambridge, MA: Cambridge University Press.

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

Machery, E. (2009). *Doing without Concepts*. Oxford, UK: Oxford University Press.

Machery, E. (2011). Replies to Lombrozo, Piccinini, and Poirier and Beaulac. *Dialogue*, 50(1), 195-212.

Machery, E. (2014). Concepts: Investigating the Heterogeneity Hypothesis, in Justin Sytsma (ed.), *Advances in Experimental Philosophy of Mind*. London, UK: Bloomsbury Publishing.

Machery, E., & Seppälä, S. (2010). Against hybrid theories of concepts. *Anthropology & Philosophy*, 10, 97–125.

McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.

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

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

Osherson, D.N. & Smith E.E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9(1), 35-58.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: The MIT Press .

Piccinini, G. (2011). Two Kinds of Concept: Implicit and Explicit, *Dialogue*, 50(1), 179-193.

Raftopoulos, A. & Müller, V. (2006). The nonconceptual content of experience. *Mind and Language*, 27 (2), 187-219.

Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.

Spelke, E.S. (1994). Initial knowledge: six suggestions. *Cognition*, 50, 431-445.

Spelke, E.S. & Kinzler, K.D. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96.

Stanovich, K. & West, R. (2000). Individual differences in reasoning: Implications for the rationality debate?. *The Behavioural and Brain Sciences* 23, 5: 645- 65.

Tversky, A. & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90 (4): 293–315.

Tye, M. (2006). Nonconceptual content, richness, and fineness of grain. In T. Szabo-Gendler & J. Hawthorne (Eds.), *Perceptual Experience*. Oxford, UK: Oxford University Press.