

views of their nature. It may be premature, however, to give up such views completely. Here I defend the possibility of hybrid models of concept structure.

In *Doing without Concepts*, Machery (2009) provides a service to us all by reminding us of the challenges of specifying what concepts are and how they are mentally represented. Moreover, by moving to the radical position that we should do away with the concepts altogether, he forces all of us to think more deeply about why we might want to preserve such notions. Finally, he is correct in pointing out the ways in which philosophical and psychological approaches to concepts often seem to be asking and answering different questions. Despite all this, it seems too extreme to assume that more traditional notions of concepts are bankrupt or that philosophers and psychologists are always talking past each other.

Here I focus on one alternative to Machery's proposal that he seems to dismiss too lightly – the idea that concepts might have a hybrid structure. I am not yet sure about the extent to which the potentially hybrid facets to concepts are actually parts of the concepts per se, or whether they should instead be considered as linked to concepts that are themselves much simpler atoms in the manner described by Fodor (1998). Cognitive science still has not fully answered Fodor's reasons for doubting that many cognitive phenomena associated with concepts reflect internal structural features of concepts as opposed to aspects of how we use and work with concepts; but to accept Fodor's arguments is to discount Machery's view as well, so let us assume here that we do want to assume internal structures to representations of concepts and that those structures help explain many psychological phenomena associated with concepts such as induction, categorization, and conceptual change.

Fodor (1998) has characterized concepts as “the smallest units of thought” and, in this respect, many psychologists and philosophers agree. Does such a characterization compel us to Machery's heterogeneity hypothesis, namely, that we must have several distinct concepts of water because, depending on context, we seem to use the concept in different ways? It is difficult to see why. Machery discounts the hybrid alternative by arguing that people will endorse conflicting statements about kinds such as tomatoes, whales, and the like, and those conflicts can only be explained by assuming that they are drawing on different concepts. These different uses are supposedly not “coordinated,” and therefore people cannot be referring to the same concepts. But this coordination problem does not seem to be so lethal for hybrid views. If a given concept has a hybrid structure consisting, for example, of typicality-based information, causal schema, functional relations, and logical entailments, it might well be the case that different contexts cause people to weight those properties very differently and respond in different manners across tasks. Ever since Lakoff's (1972) demonstrations that different “hedgies” such as “technically speaking” and “loosely speaking” can cause us to categorize kinds such as whales differently, it has been known that some hallmark ways we use concepts, such as categorization, can show strong variations as a function of situational and sentential concepts. But, if hybrid models are right, they seem more than adequate for dealing with such phenomena. Machery would need to provide a detailed model of internal hybrid representations of concepts that showed how they were intrinsically unable to computationally model such effects, and he has not yet done so.

Machery suggests that the parts of hybrid concepts must be “coordinated” such that this cannot create inconsistencies, such as categorization judgments that whales both are and are not fish. This coordination property is seen as an essential part of hybrid models, and hybrid models are described as incoherent or empty without it. This was not an obvious conclusion. Consider, for example, contexts in which we might describe a person as “short” and then “tall.” If we see a 2-meter person practicing with Olympic gymnasts, we might well call him tall; but when observing him practice with an Olympic basketball team,

we would call him short. We may well know his true height and the true heights of the other players, but the contexts call for different ways of assigning thresholds on the vertical dimension that we would then use to consider someone tall or short. Is this to be taken as evidence for multiple concepts of tall and short? Are there then an indefinitely large number of such concepts that are depending on all the micro-contexts that could shift the thresholds to tiny degrees? There is a strong tendency to resist such a route, and it seems that, for similar reasons, we should resist claims that hybrid structures are undermined by conflicting categorization judgments in different contexts. Categorization inconsistencies do not pose a problem if there are still systematic ways that categorization judgments can be shown to vary across contexts as a function of a description of their internal hybrid structure. Machery would be correct in pointing out that such systematic accounts are not yet fully worked out, but there are no obvious reasons why they might not be in the long run.

Hybrid approaches also have other appeals. They can, for example provide continuity and coherence to models of conceptual change over time, as, for example, when the causal or rule-based aspects of a concept become more differentiated as a child grows older and come to be weighted more and more relative to the associative components (Keil 1989; Keil & Newman 2010). Machery's heterogeneity alternative sees the child as progressing through a series of unconnected concepts that somehow magically tend to unfold in the same way across children. Hybrid models can also help explain how concepts differ across broad categories such as natural kinds and artifacts, where different components of the hybrid might be present to different degrees and accordingly assigned different weights, as well as being processed in different ways (Hampton et al. 2009).

Machery has done us all a great service. He raises a host of interesting troubles for many accounts of concepts, and he is to be commended for trying to build a larger common ground of inquiry between philosophers and psychologists. His book is a refreshing new perspective that prods all of us to further develop our own theories of concepts.

The function and representation of concepts

doi:10.1017/S0140525X10000397

Sangeet S. Khemlani^a and Geoffrey Goodwin^b

^aDepartment of Psychology, Princeton University, Princeton, NJ 08540;
^bDepartment of Psychology, University of Pennsylvania, Philadelphia, PA 19104.

khemlani@princeton.edu

ggoodwin@psych.upenn.edu

<http://www.princeton.edu/~khemlani>

<http://www.psych.upenn.edu/people/ggoodwin>

Abstract: Machery has usefully organized the vast heterogeneity in conceptual representation. However, we believe his argument is too narrow in tacitly assuming that concepts are comprised of only prototypes, exemplars, and theories, and also that its eliminative aspect is too strong. We examine two exceptions to Machery's representational taxonomy before considering whether doing without concepts is a good idea.

In *Doing without Concepts* (Machery 2009; hereafter *DwC*), Machery proposes that heterogeneity in the mental representation of “concepts” is sufficient to render that term useless. As he argues, the term can refer to exemplars, prototypes, and theories. However, it can also refer to defaults (Connolly et al., Fodor et al. 2007), aspects (Prasada & Dillingham 2009), Boolean concepts (Feldman 2000; Goodwin & Johnson-Laird 2010; submitted; Shepard et al. 1961), and connections yet to be discovered. Thus, in our view, Machery's taxonomy is too narrow, and it underestimates the degree of heterogeneity that

exists in the representation of concepts. It excludes a variety of conceptual phenomena that do not fall within the purview of prototypes, exemplars, and theories. We turn next to describe two examples of such conceptual phenomena – Boolean concepts, and connections expressed by generics.

Boolean concepts – those that are composed out of negation (*not*), disjunction (*or*), and conjunction (*and*) – are an important kind of everyday concept. They occur frequently in the form of laws, rules, or procedures. Indeed, Machery's criteria for individuating concepts are themselves Boolean concepts (see sect. 2 of Machery's Précis of *DwC* in this issue). A concept does not need to be entirely Boolean in order for it to contain relevant Boolean structure, however. How individuals learn Boolean concepts is still not resolved, but many of the current leading contenders are not based on exemplars, prototypes, or theories (Feldman 2006; Vigo 2009). We have recently proposed an alternative theory, based on mental models, which analyzes the complexity of concepts in terms of the number of distinct possibilities that a concept can be compressed into (Goodwin & Johnson-Laird, submitted). This theory predicts the acquisition of Boolean concepts as well as, if not better than, the other leading contenders, and it too is not based on the representational mechanisms that Machery assumes to be exhaustive in explaining conceptual knowledge.

Recently, we published a paper documenting the occurrence of "conceptual illusions," in which people think that particular instances of a Boolean concept are possible when in fact they are not, and vice versa (Goodwin & Johnson-Laird 2010). A typical example is the concept: *green and large, or else green*, in a context in which all possible objects are either green or not, and either large or small (the "or else" here represents exclusive disjunction). A large percentage of individuals think that an object that is both green and large is possible given this description, which is in fact erroneous. The exclusive disjunction between the two clauses means that the only possible object is one that is both green and small. This error, as well as others like it, is predicted by the mental model theory's principle of conceptual truth. And as far as we can tell, accounts based on prototypes, exemplars, or theories have no way to explain these errors.

Other examples of conceptual phenomena unaccounted for by the taxonomy described in *DwC* include the connections and relations that link concepts together. Such connections can be concepts unto themselves, and are revealed by generic assertions such as "tigers are striped," "barns are red," or "ticks carry Lyme disease," which express generalizations about kinds of things (Carlson & Pelletier 1995; Gelman 2003; Lawler 1973). All three assertions are true for different reasons, and as such, generics provide a means for studying the types of connections we represent between concepts of kinds and properties. We have found that for statements such as "tigers are striped," the relation between the kind ("tigers") and the predicate ("are striped") can be distinguished from logical, statistical, and causal connections (Khemlani et al., submitted; Prasada et al., submitted). These distinctions could account for phenomena in concept learning and conceptual development without importing any assumptions made by other theories of concepts. Thus, by studying generics, it is possible to discover the conceptual structure of generalizations without assuming the representational structure of the concepts to which they refer. Machery proposes that future research should examine differences between generics, and particularly how they differ from quantified assertions (*DwC*, p. 200), and we agree wholeheartedly. Unfortunately, the proposals in *DwC* do not leave room to explore such advances in conceptual organization, as they encourage researchers to couch their work as falling within the domain of three fundamental classes of conceptual representation.

We do not think these phenomena, which point to even greater heterogeneity in the mental representation of concepts than Machery suggests, strengthen Machery's eliminativist argument to do away with concepts. The elimination of the term

"concept" in favor of greater specificity tacitly endorses the assumption that prototypes, exemplars, and theories (and whatever else), are all that comprise concepts. It thus presupposes that contingent facts about the mental representation of concepts are the sole criterion for deciding whether "concept" ought to be preserved. But this presupposition ignores the common *function* that diverse sorts of concept play in representing knowledge and in communication. Concepts represent and convey systematic bodies of information, and they would retain this function regardless of how they are mentally represented. In other words, we think that the question of what counts as a concept needs answering at the computational level, not at the algorithmic one (cf. Marr 1982).

Thus, Machery's eliminativist argument is too powerful. It gives no grounds for thinking that the term "concept" is in an especially precarious position. In much the same way that the term organizes a wide array of representational processes, so too do terms like "thinking," "attention," and "memory." Would Machery have us do away with these terms as well, given heterogeneity in the cognitive processes to which they refer? Perhaps, but we think this is going too far.

In sum, the key functions of concepts are to represent and communicate knowledge, and this general functional property is what argues in favor of preserving the term "concept." We believe that heterogeneity at the level of mental representation is no obstacle to the further empirical investigation of concepts.

Concepts are a functional kind

doi:10.1017/S0140525X10000403

Elisabetta Lalumera

Dipartimento di Psicologia, Università di Milano-Bicocca, 20126 Milano, Italy, and Cogito Research Centre in Philosophy, Università di Bologna, 40100 Bologna, Italy.

elisabetta.lalumera@unimib.it

<http://sites.google.com/site/elisabettalalumera/>

Abstract: This commentary focuses on Machery's eliminativist claim, that "concept" ought to be eliminated from the theoretical vocabulary of psychology because it fails to denote a natural kind. I argue for the more traditional view that concepts are a functional kind, which provides the simplest account of the empirical evidence discussed by Machery.

The novelty of Machery's proposal in *Doing without Concepts* (Machery 2009) is the claim that the term "concept" ought to be eliminated from the theoretical vocabulary of psychology, because it fails to denote a natural kind. I will not dispute the claim that concepts are not *a* natural kind. There is a growing consensus among psychologists that the structure of concepts may vary along many dimensions, depending on expertise, domain of objects categorized, and conceptual task involved. Much of this evidence is reviewed by Machery himself, as well as in other recent studies in the philosophy of psychology (Piccini & Scott 2006; Weiskopf 2009b). My point here is that, on philosophical grounds, this evidence is perfectly compatible with the much less revisionary claim that concepts are a functional kind. Something is a concept by virtue of the function it performs within a cognitive system, and something is the concept of a certain category *C* (at least partially) by virtue of the further specific function of representing it. It is a further question whether or not the functional kind "concept" is realized by natural kinds (Weiskopf 2009b). Functional kinds can be individuated and described independently of their realizers. This, however, does not deprive them of a central role in experimental psychology.

1. Concepts are a functional kind. It is disputable that a full characterization of the psychological usage of "concepts" is: