

The plurality of concepts

Abstract: Traditionally, theories of concepts in psychology assume that concepts are a single, uniform kind of mental representation. But no single kind of representation can explain all of the empirical data for which concepts are responsible. I argue that the assumption that concepts are uniformly the same kind of mental structure is responsible for these theories' shortcomings, and outline a pluralist theory of concepts that rejects this assumption. On pluralism, concepts should be thought of as being constituted by multiple representational kinds, with the particular kind of concept used on an occasion being determined by the context. I argue that endorsing pluralism does not lead to eliminativism about concepts as an object of scientific interest.

1. Introduction

Theories of concepts often center around claims of the form 'Concepts are *K*', where '*K*' is replaced with a description of a kind of psychological structure. Claims of this form are common enough that most standard literature reviews are organized around the various kinds of representations that have been proposed to fill in the blank (Fodor, 1994, 1998; Komatsu, 1992; Laurence & Margolis, 1999; Prinz, 2002; Smith & Medin, 1981). For instance, concepts have been identified with prototypes, bundles of exemplars, theory-like structures of some sort, perceptual 'proxytypes', and unstructured atomic symbols. The debate over concepts has been focused on which of these best satisfies all of the empirical and methodological constraints on a theory of concepts.

Unfortunately, none of these theories has so far developed a clear and widely convincing lead over its competitors. The result has been a sort of impasse. No general view about what kind of representations concepts are seems entirely satisfactory. One goal of this paper is to diagnose and overcome this impasse. One possible cause is that most parties to the debate have, albeit tacitly, adopted the *uniformity assumption*: the claim that all concepts are a single, uniform kind of psychological entity. I will argue that, in fact, concepts are not constituted by a single

psychological kind, and hence that the uniformity assumption is false. It is, I claim, primarily adherence to this assumption that is responsible for the stalemate facing theories of concepts.

A second goal of this paper is to present a pluralist view of concepts, on which many different kinds of psychological structure are available to serve as concepts. Further, on this view there is no such unique thing as *the* concept of, for instance, object, living thing, tree, or elm. Many different psychological structures serve as concepts of those entities. This is true both across and within individuals: people typically have multiple concepts, belonging to different representational kinds, corresponding to any category that they are able to represent. I develop this pluralist model in some detail and defend it against the objection that it renders concepts so disunified as to be an unfit object of study for scientific psychology. I also distinguish pluralism from various hybrid models of concepts that have been proposed, and argue that hybrids are unlikely in the long run to prove successful.

2. The case for pluralism

The core argument for pluralism is as follows:

(1) Psychological theories posit that concepts play a certain causal/explanatory role in cognition.

∴ (2) Concepts should be identified with whatever entities or structures best fill that role.

(3) Several distinct kinds of psychological structures satisfy the causal/explanatory role of concepts.

∴ (4) Concepts are constituted by several distinct kinds of psychological structures.

The general form of the argument is familiar. A similar argument underwrites functionalists' identification of mental states with the physical properties or mechanisms that realize them. Functionalists maintain that mental states are members of a certain functional kind. If we want to know what mental states are, we can answer this question in two ways: either by spelling out the functional role that defines those things, or by specifying the structures that realize that functional role in a particular kind of organism. Questions concerning functional kinds can be answered either at the level of roles or realizers.

There has been extensive philosophical discussion of the notion of a kind (and the more restricted notion of a natural kind), certainly too much discussion to summarize here (see, e.g., Boyd, 1991; Hacking, 1991; Millikan, 1999). But there is little consensus on the proper analysis of kindhood. Therefore, rather than commit to a particular analysis of what makes something a kind, I will work with a more general notion of kindhood that is widely accepted. Kinds as understood here are groupings of entities that participate in a body of empirically discovered reliable generalizations, and which participate in those generalizations due to some set of properties they have in common. I leave it open just how these generalizations are discovered, how reliable they need to be, and what the relevant underlying properties might be (e.g., internal structure vs. common etiology).

We can proceed to characterize concepts, as a functionally specified category, at the level of role or realizer. Once we specify their explanatory or functional role, concepts *just are* whatever entities perform the relevant cognitive tasks. I suggest that concepts should be thought of as (1) mental representations that (2) are employed in categorization and (3) are capable of

combining productively and systematically into larger and more complex structures.¹ I will elaborate briefly on each condition.

Representation: Concepts refer to, pick out, denote, or otherwise stand in an intentional relation to individuals or sets of entities, events, states, properties, etc.² A core part of how we individuate concepts is by their representational properties—in particular, by what they are concepts *of*.³ There may be other respects in which we can also individuate them; for example, by *how* they represent what they do, or by properties they possess as representations that are not content properties (e.g., their formal or syntactic properties). When I identify concepts here, it will be primarily by the content that they represent and by how they represent it; a CAT concept *is* that concept because it represents cats in whatever way that it does.

Categorization: Categorization is often interpreted narrowly as the cognitive act of inferring membership in C on the basis of membership in various cue categories F_1, F_2, \dots, F_n : it's got feathers, it's in a tree, it's chirping, hence it's a bird. On the other hand, contrast this with inductive and deductive inference, which involve inferring membership in various target categories F_1, F_2, \dots, F_n on the basis of known membership in C : it's a bird, hence it flies, builds a nest, and lays eggs. At a formal level, categorization and inference seem to be similar. Interpreted in a narrow sense, categorization can be seen as an act of *inferring* from multiple categories to membership in a single category; similarly, induction and deduction can be seen as narrow *categorization* of an object under multiple categories, given its membership in a single cue category. Given the similarities in their logical structure, I will call both sorts of inferential

¹ Adopting this functional/explanatory role is justified by the fact that it captures much of what psychological theorists have been aiming at in proposing particular accounts of concepts. For evidence of this fact, note that these three conditions are included on virtually all lists of desiderata for a theory of concepts. Nor does it seem likely that narrowing the role by adding further restrictions will affect the core of the argument made here.

² Categories are distinct from concepts. Categories are what concepts represent; they are what our concepts are *of*.

transitions ‘categorization’. So concepts are psychological entities that play a central causal role in mediating inferential transitions and reasoning processes of this sort.

Combination: Concepts can combine productively and systematically to form larger representational structures (Fodor & Pylyshyn, 1988). We can represent not merely cats, but black cats, angry black cats, declawed angry black cats, declawed angry black cats with one white paw, and so on. This illustrates the productivity of the conceptual system: there is a potentially open-ended range of conceptualized thoughts that we can entertain. Systematicity is illustrated by the fact that the ability to entertain one thought, e.g., that some clowns like firemen, carries with it the ability to entertain related thoughts, e.g., that some firemen like clowns. The ability to grasp a single proposition from a systematically related set goes along with the ability to entertain the rest of the propositions in the set. Productivity and systematicity are among the key pieces of evidence for the claim that thoughts are structured mental states having concepts as their constituents. The combinatorial system as a whole must be sufficient to generate the entire range of possible complex representations on the basis of the elementary concepts plus whatever operations of composition are available.

With this sketch of the function of concepts in hand, we can turn to the core argument for pluralism. Premise (1) makes the claim that concepts are a certain sort of theoretical entity.⁴ In cognitive psychology, we find concepts being ascribed a central role in explaining how we carry out certain sorts of tasks and being implicated in a range of different processes. What will make a type of psychological structure count as a concept, then, is determined by the role that concepts are supposed to play in our well-developed theories of cognition. This is just to say, then, that

³ The details of how representation works—what makes something represent what it does, for example—is a question for a theory of content. I will have nothing to say in this discussion about how concepts *acquire* their representational properties. For one view of how to understand conceptual content, see Weiskopf (forthcoming *a*).

concepts are like many other entities posited in psychological theories: they are a functional kind. From this fact follows premise (2), which claims that insofar as concepts are a functional kind, they should be empirically investigated in the same way that other functional kinds are. Empirical investigation sheds light on the inner nature of functional kinds by uncovering the structures and mechanisms that enable them to function as they do. So in the case of concepts, the task is to discover precisely which structures are responsible for the various functions and effects delimited by psychological theories of concepts. Finally, premise (3) of the core argument is straightforwardly empirical. The case for (3) turns on what sorts of structures actually fill the concept role in human cognition. Given the functional specification of concepts, what things (if any) satisfy that specification?

Here there are only three possibilities: (1) *One* kind of entity satisfies the role in question, (2) a *variety* of different kinds of entities satisfy the role, or (3) *no* entities satisfy the role. In the first case, the identification question is readily answered: the functional kind can be identified with the unique kind of mechanism that realizes the role. In the last case, the question is also easy to answer: the functional specification does not pick out anything; in fact, the kind does not really exist at all. The issue is what to say about the intermediate case in which we have a variety of realizers. In order to spell out what this intermediate case is committed to, it will help to chart some of the logical geography of theories of concepts.

Theories of concepts can be distinguished in terms of how they stand with respect to two distinct assumptions. The first of these I call the *Singularity Assumption*:

⁴ That is, concepts are entities that play a central role in psychological theories; this isn't to say that they are *themselves* theories (as in the view of Murphy and Medin, 1985, for example).

(SA): For any category that can be conceptually represented, there is such a thing as the unique concept of that category.⁵

To endorse SA is to assert that we can sensibly refer to ‘*the* concept of C’, for more or less any noun phrase C. The definite article is important here, since it does the work of ruling out the possibility of there being *more than one* thing picked out by the expression ‘concept of C’. Even if there are many ways to represent C in thought, at most one of those C-representations can count as the concept of C. So, for example, anyone who represents lemurs, but doesn’t represent them in the right sort of way simply doesn’t have LEMUR,⁶ since having that concept requires having a representation of the right sort (leaving open the question of what other constraints might need to be satisfied in this case).

Proponents of SA have been most common among philosophers who favor conceptual analysis. The goal of conceptual analysis is to produce a certain kind of definition for a concept, but not any extensionally adequate definition is an analysis. Indeed, not even any definition that is extensionally adequate and *modally* adequate will do. An analysis is a definition that captures some essential and *explanatorily significant* properties had by things falling under the concept. On this view, even if there are many ways of thinking of penguins, only the one that constitutes an analysis is genuinely the concept PENGUIN.

On the side of rejecting SA are psychologists, anthropologists, and linguists who talk about a certain culture’s concept of living things, a subgroup or community’s concept of living things, the child’s concept of living things, and so on. On all of these views, concepts such as

⁵ The reference to *conceptual* representation is to rule out any non-conceptual ways that a person might have of representing a category. We could make SA more precise by adding a reference to the subject population that is capable of representing the category in question. I will assume that we are restricting attention here to human concepts (although I do not think that humans are the only concept possessors).

⁶ Small caps are used to mark the names of concepts.

LIVING THING are always *indexed* to an individual, group, or culture. Believers in SA would have to say either that one of these concepts is genuinely *the* concept of living things, or that none of them should be thought of as concepts at all (but perhaps as something else, say a person or a culture's *conception* of living things).

The second assumption concerns whether all concepts are members of the same psychological kind. I have called this the *Uniformity Assumption*:

(UA): All concepts belong to a single psychological kind.

UA should be read as claiming that concepts can be reductively identified with a particular kind of psychological structure (a type of mental representation). For instance, in classical conceptual analysis, they can be identified with a certain kind of definition. Theories that deny UA claim, by contrast, that concepts in general fall into several distinct psychological kinds. They cannot be reductively identified with any single kind of representational structure.

Evidence that UA is widely adhered to comes from looking at how arguments among theories of concepts are decided. Consider the (now-discredited) view that concepts are mental images. A standard argument against this claim is to show that many things can be *thought about* readily, but can't be *imaged*, either because they are abstract, and hence don't have any physically perceivable instantiations (e.g., JUSTICE) or because they are perceptually heterogeneous, and hence subsume a variety of more specific perceivable instantiations (e.g., FURNITURE). So not all concepts can be images. The form of argument here is straightforward: demonstrate some phenomenon that a single kind of representation cannot account for, and you undermine the theory that gives that representational kind a special pride of place. The history of theories of concepts is characterized by shifts from privileging one psychological kind to

privileging another. But the motivation for such shifts disappears if we remove the governing assumption that all concepts must belong to a single representational kind.

Views that combine the Singularity and Uniformity Assumptions I will call *monolithic* theories.⁷ Monolithic theories maintain that there is such a thing as the unique concept MOTHER, for example, and also maintain that all concepts (MOTHER, ELM, CHASE, WHISKEY, CATCH, etc.) are members of the same psychological kind. This combination of views has probably been the norm among theorists of concepts. I have already mentioned the view of concepts tacitly held by conceptual analysts as one instance of a monolithic position. Consider also traditional prototype theorists, who held that all concepts are structured like prototypes, and that having the appropriate prototype for a category just *is* possessing the concept of that category (Hampton, 1995).

The Singularity and Uniformity Assumptions cross-cut one another, so it is coherent to hold each pairing of SA vs. non-SA with UA vs. non-UA. One could combine non-SA and UA: on this view, there is no such thing as the unique concept MOTHER (that is, people possess many distinct MOTHER concepts), but nevertheless all concepts belong to a single overarching psychological kind. An illustration of this would be a theory on which all concepts are prototypes, but are distinguished from one another by slightly different features and weights. So I might possess a MOTHER concept that assigns degree of biological relatedness a higher weight than does yours, on which nurturing and care get a high weight. These are both MOTHER prototypes, but in virtue of their different weighting of these features they are *different* prototypes, and hence different concepts.

⁷ Lance Rips (1995) also uses the term ‘monolithic’ in passing to describe theories of concepts that posit only a single kind of representational structure.

Alternatively, one might combine SA and non-UA. On this view, there would be a unique MOTHER concept, but it might not belong to the same psychological kind as, say, JUSTICE, WATER, or BEAR. On such a theory, perhaps abstract, substance, and natural kind concepts are structured differently from kinship concepts; but nevertheless, for each category that can be represented, there is a single representation that serves as *the* concept of that category. Something like this view is advanced by domain theorists, who hold that concepts that represent categories belonging to different ‘ontological domains’ are structured differently; for instance, natural kind concepts are distinct from artifact concepts, concepts of animate creatures are distinct from concepts of substances or stuffs, and so on (Hirschfeld & Gelman, 1994; Gelman, 2003; Medin & Atran, 1999).

Finally, it could be both that there is no such thing as the unique concept MOTHER, *and* that concepts generally are not a single kind of entity. This position, constituted by the negation of the assumptions behind monolithic theories, is what I have called conceptual pluralism. Arguing for pluralism first requires demonstrating that concepts are in fact constituted by a variety of distinct cognitive kinds. I now turn to some of the evidence that supports this claim.⁸

3. Kinds of concepts

The received view in cognitive psychology for at least a decade—from the publication of Rosch and her collaborators’ seminal papers in the mid-1970’s until the mid-1980’s—was that most lexical concepts are prototypes: representations of the statistically central properties of a

⁸ Space considerations obviously preclude the kind of thoroughgoing review of the empirical phenomena that would be needed to provide the strongest support for pluralism. The findings presented here should be thought of as a small but representative sample drawn from a much larger body of evidence.

category (Rosch, 1978; Rosch & Mervis, 1975).⁹ Combined with this is an account of categorization: something is judged to fall under a concept if it is represented as sufficiently similar to the prototype. Finally, prototype theory also offers an account of concept acquisition: concepts are learned by abstracting common features from similar encountered objects.

Prototype theory serves as an illustration of how theorists of concepts have characterized the psychological kinds that they are investigating. Concepts can be divided into kinds on the basis of:

- (1) what category they represent;
- (2) how they represent that category;
- (3) how they are usually processed; and
- (4) how they are usually acquired.

On classical prototype theory, all (or almost all) categories are represented by prototypes, which encode statistical information about what properties tend to be exhibited by category members. These prototypes are typically manipulated by processes of similarity computation over their weighted features, and are acquired by abstraction over instances. Generalizing this account of representational kinds, we can say that theories are in competition when they posit representations that differ in any of these four ways. Various competing theories have posited representations that belong to different kinds (by these criteria) to account for the experimental results. The question is whether any one of these theories can account for all, or even most, of the data.

⁹ It should be emphasized that Rosch herself did not believe that prototypes were mental representations of any sort. At many points she seems opposed to the very notion of mental representation, in fact. However, it is nevertheless true that the dominant reception of Rosch's results has involved interpreting them in the framework of representations, processes, systems, mechanisms, resource stores, and other cognitivist apparatus (see, e.g., Hampton, 1995). So while it's strictly inaccurate to ascribe to her belief in the prototype theory in the sense laid out here, her work is nevertheless naturally taken to support this view.

One early challenge to prototype theory came from exemplar theorists (Estes, 1994; Medin & Schaffer, 1978; Nosofsky, 1986, 1988, 1991). Exemplars are representations of individual category instances stored in memory. Categorization, typicality judgments, and other inferences are explained by the degree to which a target object is similar to the exemplars one has stored. Exemplar theory thus employs similarity computations, as does prototype theory, but posits a host of detailed, particular representations rather than a single summary representation of a category. On this view, there is no need to abstract and store such a summary, since the individual exemplars themselves can explain people's performance.¹⁰

Despite their virtues, the limits of prototype and exemplar theories have become increasingly obvious in recent years.¹¹ Since their problems are related, I will discuss only prototype theory here. I focus on two phenomena that demonstrate how the theory falls short of capturing all of the ways we conceive of and reason about the world: the existence of concepts structured around category ideals, and the use of causal models in categorization.

To many, similarity-based processing and prototype representations are synonymous. But they can be dissociated: not all similarity-based processing involves prototypes. Barsalou (1983, 1985) showed that many concepts are organized around similarity to *ideals* rather than prototypes. A category ideal is a representation of the best features or best exemplars of the category—often the ones that display properties important to the purpose of the category. These ideals may never actually be instantiated; for example, the ideal diet foods might be those with zero calories, although no *actual* diet foods have zero calories, and the *prototypical* diet food has

¹⁰ Exemplar theories are exceptions to the generalization that traditional theories of concepts have been committed to the Singularity Assumption. The whole point of exemplar models is that there are no single summary representations stored in memory; concepts are thought of as bundles of representations of particular instances. Not all exemplar theorists reject summary representations, but they emphasize that such singular concepts cannot be the whole story about categorization.

¹¹ On empirical successes of prototype theory, see Murphy (2002, chap. 2); for a review of exemplar theory, see Murphy (2002, chap. 4).

more than zero calories (Lakoff, 1987). Ideals have sometimes been assimilated to prototypes, but this glosses over the important differences in what properties they represent. Prototypes, as I characterize them here, represent *statistically significant* properties, while ideals present the optimal or ‘best’ properties for a category. This difference in the information about a category encoded by these distinct representations suggests that they should be treated as separate kinds.

Reasoning about some sorts of things tends to promote the use of ideals. For ad hoc concepts (WAYS TO MAKE FRIENDS) and goal-derived concepts (CAMPING EQUIPMENT), goodness-of-exemplar judgments are overwhelmingly well-predicted by similarity to category ideals, while prototypes play virtually no role. In particular, closeness to the *ideal* value (rather than the typical value) on ideal dimensions matters. Natural kind concepts (APPLE, BANANA) are frequently represented as prototypes, but even in this domain, ideals account for some of the variance in people’s judgments. When a particular goal (e.g., achieving health) is salient in people’s minds, they are more likely to represent objects (e.g., foodstuffs) in terms of how well they contribute to satisfying that goal (Ratneshwar, Barsalou, Pechmann, & Moore, 2001). The existence of ideal-centered reasoning shows that prototypes are not invariably the cognitive tool of choice for representing a category. People may sometimes represent a natural kind by its prototype, sometimes by its ideal.

In some cultures and communities, natural kind concepts may even be predominantly governed by ideals rather than prototypes. Goodness-of-exemplar judgments from North American bird experts and members of the Itza Maya community show that both groups tend *not* to make such judgments on the basis of prototypicality (Bailenson, Shum, Atran, Medin, & Coley, 2002).¹² In judging exemplar goodness and justifying inductive inferences, non-experts

¹² The Itza are an American Indian people who inhabit the rain forests of Guatemala. For further details on the Itza, see Bailenson et al. (2002).

tended to be influenced by prototypicality, while the American experts and the Itza offer justifications based on evolution, diversity and range, and other ecological factors. One proposal to explain these results is that the differences in first-hand experience with these species tend to promote representing them in terms of ecologically important (but non-obvious and non-prototypical) properties.

Experts behave like members of a separate culture in domains besides birds, as well. Fish experts categorize in terms of both overt morphological properties and non-obvious properties related to fish behavior, ecology, and usefulness to humans (e.g., whether some fish are edible or are 'junk' fish), while novices group them solely by morphology (Boster & Johnson, 1989). Burnett, Medin, Ross, and Blok (2005) found that among North American Menominee Indians culturally significant fish like sturgeon and trout are rated more typical than they are by European American fishermen from the same region of Wisconsin. Overall, combining the Menominee and European American fishermen, typicality is best predicted by how desirable and familiar the fish is, rather than by its folk or scientific central tendency (prototypicality). Tree experts judge exemplar goodness on the basis of fit with two ideals, namely whether a tree is less 'weedy' and more tall (Lynch, Coley, & Medin, 2000). Taller, less weedy trees are better exemplars of TREE. While novices tend to judge the typicality of trees on the basis of statistical central tendency, experts rely more on subjective familiarity, or breadth of individually acquired exemplar knowledge concerning trees (Johnson, 2001). Where people have domain-specific knowledge, use of prototypes may be deemphasized in favor of ideals structured around customary task demands and interactions with category members.

Ideals differ from prototypes along at least two dimensions. First, ideals contain information about the features had by the best members of a category. What makes them the best

may vary somewhat from case to case. It may be determined by aspects of the task or the culture. But the best category members are not, generally speaking, the statistically prominent category members. So ideals encode different information about categories than do prototypes. Ideals are also typically acquired in different ways. Some ideals are derived from the goals subjects have in mind in interacting with a category. Tree experts differ in the aspects of trees they find ideal depending on what kind of expert they are: a hands-on landscaper and a botanist treat different trees as ideal (Lynch et al., 2000; Medin, Lynch, Coley, & Atran, 1997). Another possible source for category ideals is the culture at large. Ideals for gender, for instance, are probably highly conditioned by culture, but so are some biological categories (Burnett et al., 2005). So goal-based reasoning and internalization of cultural norms drive acquisition of ideals, in contrast with many prototype models' emphasis on abstraction of concepts from experience. In terms of processing, however, ideals and prototypes both rely on similarity computations in delivering categorization judgments.

A further deficiency of prototype theory is its inability to account for how categories are represented in causal and counterfactual reasoning. Even young children are willing to set aside prototypical information about an animal's appearance and behavior when they are asked to categorize counterfactual instances that depart from the perceptual prototype (Keil, 1989), a tendency adults share (Rips, 1989). When asked to consider what they would say if they encountered something that has the prototypical appearance of a zebra, but which has that appearance because it was dressed up in a costume, most children judge that this is no zebra. On the other hand, if the appearance is due to surgery or other medical treatment, they may be less sure. This illustrates that what is prototypical may be useful in identification and certain kinds of

default reasoning about a category's properties, but in modal contexts it should often be discarded.

Many judgments about counterfactual scenarios are driven not by prototypes but by causal models, which represent how people think properties belonging to category instances depend on other such properties (Rehder, 2003; Sloman, Love, & Ahn, 1998). Some properties of a category are thought of as more mutable or changeable than others, and how mutable a property is depends not on its weight in a prototype, but on how many other properties of the category depend on it. A property that many other properties depend strongly on is a *central* property for a category. Centrality is a matter of a property's depth and importance in a causal model for the category. For instance, HAVING A SEAT is central to CHAIR, since lots of chair properties depend on it (YOU CAN SIT ON IT, IT IS COMFORTABLE, IT HOLDS PEOPLE, etc.). Centrality in causal models predicts performance on a variety of different tasks, including: the difficulty of imagining an instance of the category that lacks its central properties, how surprising it would be to encounter a category member that lacked such properties, the goodness of exemplars of the category, and similarity to the category ideal (Sloman et al., 1998).

Ahn and her collaborators have shown that causal models trump prototypes in many reasoning tasks (Ahn & Dennis, 2001; Ahn & Kim, 2001; Ahn, Kim, Lassaline, & Dennis, 2000). When people have causal dependency information about a category, they will often *prefer* to categorize on the basis of overlap only at the level of the most fundamental causal properties. If people are told that Jane is depressed because of low self-esteem, Susan is depressed because she has been drinking, and Barbara is defensive because of low self-esteem, they prefer to categorize Jane and Barbara (cases with matching causes) together. Common causes can override overlap of features—even if only *one* central causal feature is shared among exemplars that

differ in several more superficial features, people may categorize them together (Ahn & Dennis, 2001). For example, Rehder and Burnett (2005) taught participants about categories in which (i) properties F_2 , F_3 , and F_4 are all caused by F_1 , and (ii) properties F_1 – F_4 occur with uniform probability throughout the category. Having learned this information, people judge that an object that has F_2 – F_4 but lacks F_1 is probably not a category member. The likelihood of positive categorization in fact generally *decreases* with the number of effect features that are present: something having many of the effects but lacking the common cause is unlikely to be a category member even though it possesses many typical category features.

In terms of information represented, causal models go beyond encoding prototypical information to include conjectured mechanisms for producing typical category features. In addition, causal models generate predictions about how category members will appear and behave in counterfactual circumstances. Categorization using causal models is frequently described as an explanatory process. If a particular instance's properties could best be explained by the operation of a certain causal mechanism, then that instance is categorized as belonging to the class picked out by the model. Murphy and Medin's (1985) example of a drunk jumping into the swimming pool at a party illustrates such reasoning: the person is classified as drunk because being drunk causally explains his behavior. This process goes beyond feature-matching and similarity computation as it is usually understood in a prototype framework.

4. A pluralist theory

Taken together, it appears that the conceptual system reliably employs at least prototypes, exemplars, ideals, and causal models. Prototypes are the cognitive tool of choice for the vulgar: naïve individuals conceptualize many domains in terms of prototypes, but other common

domains promote the use of category ideals. More knowledgeable individuals' concepts are a blend of ideals derived from theoretical information and the experience of interacting practically with the category in the world. Finally, both naive and expert individuals often conceive of domains in terms of causal models, especially in counterfactual reasoning. But there is no reason to think that these models play a role in tasks such as perceptual identification, where similarity to prototypes may dominate. The same content might be represented in one situation by one kind of concept, in a different situation by another kind. What this suggests is that the conceptual system is structured to employ a variety of different representational tools. These representations differ in the kind of information they encode, some of the ways they are acquired and processed, the domains they favor, and the tasks that promote their use. And since these properties qualify them as belonging to different cognitive kinds, we should reject UA.

This evidence can also be parlayed into an argument against SA. We have seen that people can be induced to represent and reason about the same categories in a variety of different ways. But for something to be a concept is just for it to function in the right way in cognition—for it to represent a category, be capable of combination with other representations, and be causally central in categorization. The representations described in the previous section all satisfy this specification. There is nothing that would privilege one of these over any of the others.

Others have also rejected SA on grounds similar to these. Barsalou (1987) noted that judged typicality of a category's instances was highly unstable across and within individuals, and explained this instability by differences in the particular concept of that category being entertained in each context. Prinz's proxytype theory follows Barsalou in representing concepts as temporary constructs in working memory. Proxytype theory differs from pluralism in two

respects, though. First, consistent with Prinz's empiricism, proxytypes are supposed to be *perceptual* representations, or copies thereof. But pluralism as such contains no inherent commitment to empiricism; for some doubts about the viability of concept empiricism in general, see Weiskopf (forthcoming *b*, 2007). Second, Prinz often seems to think of proxytypes themselves as being essentially prototype-like. Barsalou also did not at the time explore the idea that still-greater instability might be attributable to the presence of different *kinds* of concepts. Pluralism focuses attention on the fact that there is a much wider range of representation types available to serve as concepts than just prototypes.

To illustrate how pluralism works, consider Amira's many CAT concepts. Having encountered many particular cats, Amira has many cat exemplars stored in memory (Estes, 1994; Medin & Schaffer, 1978). As she has encountered particular cats and learned more about cats in general she has constructed a cat prototype that represents the statistical features of catkind. At the same time, she believes various things about cats as a natural biological kind, such as that what makes them cats has something to do with what's inside them and what their biological origins are, plus what kinds of transformations they can undergo. Perhaps she regards some of this knowledge as being about what is essential to cats (Gelman, 2003). She also believes that some properties of cats causally depend on others (e.g., PURRING depends on HAPPINESS, MEOWING depends on HUNGER). All of this vast body of information coexists in long term memory, but not all of it is (or could be) activated at a single time. Portions of it, though, are available for ready extraction and use in categorization, building new representations, and guiding actions. When some portion of this information is activated and tokened in working memory, it constitutes one of Amira's CAT concepts. This state of affairs is illustrated in Figure 1.




Figure 1. A pluralist model of concepts. Arrows indicate retrieval from long-term memory.

There are indications that people do in fact tend to select and retrieve from memory only those category representations that will be useful given their background, current goals, and the requirements of the task. One criticism of early research on exemplar models, in fact, claims that their success is due to the fact that the artificial categories that people are trained on are small, unstructured, and nearly impossible to learn by any strategy *other* than exemplar memorization (J. D. Smith & Minda, 2000). Using an on-line priming technique, Malt (1989) found that participants could be induced to retrieve either exemplars or prototypes, depending on the particular task demands. She found that where using memorized exemplars was possible but not required, people tended to rely on abstracted prototypes instead. Where stimuli varied in their typicality, though, people tended to split their categorization decisions between prototypes (for typical members) and exemplars (for atypical members). So the strategic choice of a representation type depends on both the task and the structure of the materials.

Whittlesea, Brooks, & Westcott (1994) also showed that whether knowledge of particulars or general category knowledge was used can depend on a range of factors. For instance, in one experiment they presented two items and asked participants to decide which one was a category member. Whether participants used exemplar or prototype information depended on how the items were presented visually: side by side to emphasize their distinctness, or in a column to emphasize their common features. The former produced an exemplar memory effect, the latter produced a prototype retrieval effect. In addition, Juslin, Jones, Olsson, & Winman (2003) showed an effect of the type of feedback given during training on later categorization performance. When participants were given high-quality feedback about training items, they tended to classify test items as if they were comparing them with abstracted summary representations, while if they were given low-quality feedback they tended to classify new items on the basis of similarity to stored exemplars. So learning conditions can influence the kind of conceptual representation used in later judgments as well.

This claim about how learning influences representation has been persuasively demonstrated by Yamauchi & Markman (1998; see Markman & Ross, 2003 for a review of further evidence). In Yamauchi & Markman's studies, participants were trained on the same categories in two different conditions: a classification task and an inference task. The classification task involved learning to apply the right label to category members, while the inference task involved being given a partial description of a category member and inferring its other properties on that basis. The nature of the representation learned was probed in a transfer task involving both new and old exemplars. Those who learned the categories by classifying instances tended to do better on later classification tasks than on inference tasks; similarly, those who learned by making feature inferences tended to do better with later inference tasks. But

moreover, those who trained on classification acquired representations that included information about specific category exemplars, particularly exceptional members. Those trained on inferences acquired prototype or central-tendency representations. So different learning conditions can produce representations that underlie slightly different performance on later tasks—although these differences shouldn't be exaggerated, since even though those who learn one sort of representation could still perform other sorts of tasks with them. Moreover, the representations learned are adaptive: they are particularly tuned to carrying out a particular sort of task involving the target category.

The sort of representational pluralism outlined here is arguably the norm in cognition, rather than the exception. Consider how our best theories of perceptual systems such as vision say that they are organized. Visual representations come in a wide variety, differing in their form, content, and characteristic mode of processing. In the study of visual object recognition, for instance, there is a distinction drawn between viewpoint-based (or *image-based*) representations of objects and viewpoint-independent (or *structural description*) representations of them (Tarr & Bülthoff, 1998; Palmeri & Gauthier, 2004).¹³ An image-based representation encodes a 'snapshot' of how that object looks from a particular viewpoint—a certain angle, certain illumination and occlusion conditions, etc. (Tarr, 2003; Poggio & Edelman, 1990). A structural description representation, by contrast, encodes a three-dimensional model of an object that presents it as built up out of simpler geometric constituents in abstraction from how they look from any particular point of view (Biederman, 1987; Marr & Nishihara, 1978). These distinct representation types are tuned to certain computational tasks, and their properties may be partially explained in terms of their being suited to carry out these tasks. Indeed, it seems

¹³ My thanks to an anonymous referee for pointing out how well the debate among models of visual object identification illustrates the present point.

unlikely at this point that either single type of object representation will wholly serve the distinct purposes of visual identification and generalization across category members. If so, it would provide another example of representational pluralism in cognition, albeit in the perceptual rather than the conceptual domain.

A similar claim has been put forward about the cognitive systems that underlie humans' abilities to deal with number and quantity. Dehaene (1997), for example, proposes that numerical cognition involves three separate representations that are partially intertranslatable: an analog number line (shared with infants and nonhuman animals), a verbal representation of number names, and a visual representation of Arabic numerals. These representations are hypothesized to be activated in distinct tasks (e.g., rapid quantity estimation vs. precise numerical comparison), and to explain various patterns of deficit arising from neurological damage. Cognitive faculties like vision and number, then, seem not to rely on single inner codes for every computational task.

What I have been arguing here is that the evidence suggests that there are a variety of different kinds of category representations that the conceptual system employs as well. This fact provides the motivation for rejecting the Uniformity Assumption. The evidence canvassed so far also suggests that these representations are created and retrieved depending on the particular sort of task that is relevant at the moment (as well as more general features of the learning situation, category structure, etc.). The conceptual system, then, appears to be *adaptive*. Tasks that require detailed information about particular category members tend to promote the use of exemplars; tasks that involve reasoning about how some category properties depend on others (e.g., counterfactual reasoning) tend to use causal models; and so on, for other types of representation. These generalizations are statistical, hence there is no guarantee that this fit will be perfect in

every case, or that for every sort of task there is a dedicated type of representation that contains information that will be ideally suited to its solution. This sort of context sensitivity is compatible with there being default representations that tend to be used when there is no overriding feature of the context that would encourage a switch. Prototypes may be such defaults for most people, although as we have seen, for experts and members of different cultures, there are categories for which ideals may be the default instead.

Monolithic theories of concepts are guided by the assumption that there will be a single sort of representation that is invariably used in all or most conceptual thought, and a relatively small set of processes that operate over these representations. As an initial hypothesis about how the conceptual system works, this is not unreasonable. But monolithic theories also predict a certain amount of inflexibility in this system. If one type of representation dominates, then it will be difficult or impossible to produce representations that encode different information or are processed differently. This in turn ought to render certain cognitive tasks that depend on using such information more difficult. To take a simple example, a common criticism of pure prototype theories is that they fail to explain how everyday counterfactual reasoning about non-prototypical category members is possible. An F that fails to display the prototypical F-properties ought to be *conceptually contradictory* on these theories, since it is a conceptual truth about Fs that they display the features of the F prototype. The move away from prototypes and towards essentialist theories of representation was motivated in large part by experiments that involve asking people to reason either about Fs that fail to conform to the F prototype, or about things that conform to the F prototype but aren't Fs. But clearly, we can reason about such objects; if we couldn't, then subjects would find the experimental materials themselves to be incomprehensible. This illustrates another virtue of a pluralist theory: in rejecting the Singularity Assumption, it treats

any particular representation that we have of a category to be, in principle, open to revision or outright rejection.

So, to summarize the pluralist position: concepts are something like the representational tools deployed by a creature's central systems to categorize and draw inferences about objects in the world. Like many other cognitive systems (vision, number), the conceptual system does its job by using a variety of different information-carrying structures. The primary difference between the conceptual system and these others is that there is significantly greater latitude about precisely which concept is going to be selected for use on a particular occasion, while these other systems remain relatively constrained in how they represent the properties in their domain. But this latitude is, as we have also seen, subject to influence by a range of contextual factors that need further experimental work to spell out in detail.

Some aspects of the position sketched here have been anticipated in the literature. In a review article entitled "Are there kinds of concepts?", Medin, Lynch, & Solomon (2000) distinguished different kinds of concepts using some of the same criteria employed here (processing, content, and what they termed 'structure', by which they mean the kind of featural organization different concepts have). They concluded, in response to the question posed by the title of their paper, that there are very likely different kinds of concepts. Medin, Lynch, & Solomon don't make this case in explicitly functionalist terms. But obviously, this answer accords with the one given here, and the argument employs similar evidence: concept-involving tasks do not seem to involve the same kinds of representations, processes, content, etc., and hence concepts many belong to different kinds distinguished along these and other dimensions.

One salient difference between the position laid out by Medin, Lynch, & Solomon and the pluralist position sketched here is that while they lay out evidence in favor against the

Uniformity Assumption, they don't discuss the Singularity Assumption. So, for example, they suggest that distinct ontological domains may be represented differently (artifacts vs. kinds), as may categories given distinctive linguistic expression (count vs. mass nouns). But they don't mention the possibility that within these broad classes the same categories may be represented in distinct ways by the same individuals at different times, depending on variations in task demands and other contextual factors. The denial of UA and SA is part of what is distinctive about the position outlined here.

A similar sort of argument to the one given here has been presented by Machery (2005), and later criticized by Piccinini & Scott (2006; see also the reply by Machery, 2006). Machery wants to argue that, in his words, concepts are not a natural kind. He and I agree on roughly what it takes for something to be a kind: kinds are sets of entities that share some large body of properties because of a common underlying mechanism. Machery holds that evidence similar to that reviewed here shows that concepts have no such common underlying properties, and hence do not qualify as a kind. One consequence that he draws from this is that "the notion of concept is ill-suited to formulate scientifically relevant generalizations about the mind" (Machery, 2005, p. 464-5). Rather, he proposes that psychologists should focus on formulating generalizations involving the specific subkinds of concepts that have been discovered (e.g., prototypes, exemplars, and mini-theories). Finally, he raises (but doesn't address in detail) the provocative question of whether, if concepts are not a scientifically legitimate kind, we should continue to talk about them at all.

Unlike Medin, Lynch, & Solomon, Machery does note that we typically possess multiple representations of the same category; that is, he endorses the denial of the Singularity Assumption. However, there are two major differences between the pluralism I advocate and

Machery's position. First, I take it that the empirical data support the view that the conceptual system as a whole is adaptive; that is, the particular system of processes and representations that are pressed into service on a particular occasion tend to be appropriately fitted to performance on the task at hand. While Machery notes that there are many such systems, he doesn't focus on the issue of how their operations are coordinated and deployed.¹⁴ Second, Machery denies that a psychological theory of concepts *per se* is in the cards. This difference between Machery's view and the one advocated here can be traced to our relative degrees of pessimism (or optimism) about the ultimate scientific utility of retaining the notion of a concept in psychology. I will argue in the next section that concept nihilism of the sort Machery advocates is unwarranted. To foreshadow a bit, I suggest that the fact that a class of entities is partially heterogeneous at one level of analysis doesn't preclude their being homogeneous at a higher level of analysis. This is a common feature of superordinate kinds in nature, and concepts, I propose, are such a superordinate kind. So while Machery and I have a similar view of the import of much of the empirical data—we both take it to support non-Uniformity (what he calls the 'heterogeneity hypothesis') and non-Singularity, I take this to illustrate a deep and important fact about the adaptive character of our concepts, while Machery takes it to show that concepts are not an interesting or unified object of study at all.

The point about superordinate kinds is, in fact, foreshadowed by Piccinini & Scott (2006, p. 398), who note in their critique of Machery that even if the phenomena associated with conceptual thought (categorization, inference, comprehending language, productive combination, etc.) were underwritten by distinct kinds of representations, it is a further step to the conclusion

¹⁴ Although I should note that he isn't committed to denying the adaptiveness claim, either. This simply reflects our different focus. Whereas I hold that there are interesting general facts to be discovered about concepts and the conceptual system, Machery is committed to denying that concepts hold any scientific interest as such. Hence his

that concepts *as a whole* aren't a kind, since reaching this conclusion requires ruling out the possibility that concepts can be subsumed under a higher-level kind. Piccinini & Scott are primarily concerned to lay down conditions under which we can legitimately infer that conceptual representations divide into separate kinds. They make two main proposals. First, linguistic understanding and inference, which are usually taken to be among the phenomena to be accounted for by a theory of concepts, require distinctive sorts of representations. Lexicalized concepts, then, are likely to have a different form than non-lexicalized concepts, which don't depend on the specialized sorts of representations employed in linguistic processing. Second, concepts can be separated into kinds according to whether they are involved in similarity-based computations or whether they are not (as, they allege, in the case of ad hoc concepts and logical/mathematical concepts).

While, for reasons of space, I have not dealt explicitly with the lexicalized/nonlexicalized distinction here, it seems plausible that lexicalized concepts are in fact somewhat special.¹⁵ It may even be the case that lexicalized concepts behave differently when their specifically lexical properties are relevant. For example, Malt, Sloman, Gennari, Shi, & Wang (1999) found that while speakers of different languages (Chinese, Spanish, and English) sorted objects together differently when the task required doing so on the basis of their conventional linguistic label, they sorted them into essentially the same similarity clusters when they were asked to do so without considering how they are to be named. Naming, then, induces people to see the same objects as falling into different sorts of similarity clusters. This finding gives some support to the

endorsement of the adaptiveness claim would have to be contingent on reading it as something other than an interesting empirical truth about concepts as a whole.

¹⁵ This leaves it open just what the source of this specialness is. It might be that only certain sorts of concepts can become lexicalized, in which case the specialness of lexical concepts arises from the existence of this filtering mechanism. On the other hand, it might be that becoming lexicalized induces changes in how a concept is represented and processed, although there are at best weak constraints on what concepts can be lexicalized. In the

proposal that when lexicalized concepts are accessed via their lexical ‘hooks’ they may behave differently than when they are accessed in other ways.

Having said this, it isn’t clear that Piccinini & Scott’s second point can be sustained. As I’ve argued, there are many distinct kinds of ‘similarity-based’ concepts: ideals, exemplars, and prototypes are all examples. Neither do these similarity-based concepts seem to line up neatly with lexicalized concepts, as they propose in passing (pp. 405-6), since typicality effects are found in both lexicalized and non-lexicalized concepts (a point also made by Machery, 2006 in his reply). So I agree with Piccinini & Scott that we need to focus on how concepts might divide up along the lines laid down by distinct tasks, and generally that we need to pay attention to the role of task demands in determining how individuals conceptualize the world. However, I don’t agree that in general there will be a simple division of concepts into kinds according to whether they are lexicalized or not, or whether they are similarity-based or not. The evidence suggests that the same categories (even lexicalized ones) can be represented differently on different occasions, and also indicates that the same categories can be simultaneously represented, within a single individual, by many ‘similarity-based’ concepts.

5. Pluralism versus nihilism

I turn now to Machery’s contention that pluralism leads us to a kind of eliminativism, or *nihilism*, about concepts. Machery’s point, in essence, is that pluralism undermines concepts as a theoretical category of interest to psychologists. If concepts are not reducible to a single, uniform kind, but are rather a collection of distinguishable subkinds, do they have any role to play in a scientific psychology, or should they be eliminated in favor of some other construct—perhaps

latter case, lexicalization is a cause of the specialness of these concepts, rather than a consequence. I take no stand here on which of these is the right view.

talk only of the various particular kinds of structures that people employ in reasoning about categories?

Generally speaking, a nihilist challenge has the following form: what are we to do if, given a functional specification F , a number of distinguishable kinds of entities K_1, \dots, K_n individually satisfy F ? Nihilists say that in such a situation we should abandon the idea that F picks out a category of scientific interest. Perhaps we might split F into a set of distinct subkinds, perhaps we should simply get rid of F -talk altogether. In any case, though, a plurality of realizers for a single functional role puts pressure on the scientific integrity of the category picked out by that role.

This is a familiar style of argument against functional kinds (Block, 1997; Craver, 2004; Kim, 1992). We should, however, distinguish two styles of nihilism. One style is straightforward eliminativism. Eliminativists simply claim that there are no concepts at all. If *nothing* satisfied the functional role of concepts, this conclusion might be warranted. Eliminativists about mental representation, for instance, are concept nihilists by implication. This form of nihilism is not at issue here. The weaker form of nihilism is one that claims that concepts *exist*, but do not constitute a kind. But being a kind is often regarded as a prerequisite for a category's being of scientific interest, since the sciences strive to discover kinds and to map their internal structure and relations to one another. Kinds are, minimally, supposed to underwrite a substantial body of empirically discovered generalizations, so if a cognitive category fails to be a kind, it falls out of the purview of psychology. If concepts are a mere grab-bag of unrelated kinds of representations that are implicated in no higher-level generalizations, this weaker form of nihilism might be justified.

Nihilists might point to the fact that differences in the kinds of concepts we employ can also give us more fine-grained explanations and predictions than the assumption of a common structure would. If people tend to use a certain kind of representation or process under certain conditions, then we can predict how they will behave on certain tasks better than if we assumed that they always use the same kind of representations and processes. Pluralism can account for the causal differences between two instances of thinking about a category by appealing to these contextually restricted generalizations. Given this advantage, it is perhaps unclear why we should advert to the higher-level kind at all.

Nevertheless, I will argue that we should not embrace concept nihilism. To say that there is no theoretical interest in concepts proper, but only in the narrower classes of prototypes, ideals, exemplars, causal models, etc., would be to ignore the significant explanatory and functional commonalities that those entities have. To begin with, there are the functional commonalities incorporated in the concept role itself. Prototypes, causal models, etc. all serve to represent categories in thought, have the capacity to be combined into larger structures, and all play a role in categorization. One might object, though, that these properties are not empirically discovered, but instead comprise the functional role that *defines* being a concept. Hence they provide no support for the claim that concepts are a kind.

Even if we accept this, none of the evidence cited so far precludes the existence of empirical generalizations that are comparatively *insensitive* to the particular kind of concept being used. If there are such generalizations, the attraction of nihilism drops away. I maintain, then, that there is a sense in which concepts *are* a kind: these various kinds of representations belong to a common *superordinate* kind. This sort of organization is hardly uncommon in nature.

Animals are a kind, and they subsume numerous subkinds, e.g., octopuses, roundworms, dragonflies, and humans. Kinds often belong to such nested hierarchies.

To show that concepts are a high-level kind, we need to cite non-analytic generalizations that subsume these representational subkinds. In what follows, I will sketch four types of high-level processes and generalizations that govern concepts as a class.

1. Logical form. One example of such processes is inferences that depend on logical form. If mental states have internal, logical structure, then it is reasonable to suppose that there are mental processes that are sensitive to that structure, rather than to the particular concepts that are being combined in that structure; otherwise, how else are complex mental states constructed in the first place? And if there are processes that can construct these mental states, there must be further processes that make use of their structure.

For an example of a process that is sensitive to the internal logical structure of a mental state, consider the inference from SOPHIE IS A BLACK CAT to SOPHIE IS BLACK and SOPHIE IS A CAT, neither of which appears to depend on the subject representing cats in any particular way. Whether a subject is representing cats by means of a category ideal, an exemplar, a causal model, or some other kind of concept, these inferences obtain. This suggests that formal inference processes generalize over these different types of concepts. Consider any two occasions on which a person is thinking that Sophie is a black cat. On one occasion her thought contains her concept CAT_1 , on another it contains CAT_2 (where these two concepts belong to different kinds). But on each occasion, she can infer that Sophie is a cat. We could explain this by positing two separate inferential processes, one governing the behavior of concepts of type CAT_1 and the other governing concepts of type CAT_2 . Or we could adopt a higher-level generalization and subsume both inferences under a single process that is sensitive to the fact that each concept is playing the

same syntactic role in these distinct thoughts. That is, whenever she entertains a thought of the form *x is an ADJ N*, she can infer thoughts of the form *x is ADJ* and *x is an N*, no matter what kind of concept the *N* in question is. The existence of such formally specified inferences indicates that some mental processes are insensitive to differences in kind of conceptual structure. These are mechanisms that operate over concepts *tout court*.

2. *Conceptual combination*. How concepts are combined provides another example of processes that tap multiple kinds of concepts. In many cases, conceptual combination results in ‘emergent’ features—features that are part of the complex concept, but not part of the constituent concepts that contributed to it. Examples are familiar from the literature: WOODEN SPOONS are typically LARGE, but neither SPOONS nor WOODEN (THINGS) are typically large; HARVARD CARPENTERS are non-materialistic, but neither HARVARD (GRADUATES) nor CARPENTERS are (Kunda, Miller, & Claire, 1990); PET FISH are small and live in bowls, which neither PETS nor FISH typically do; and so on. The existence of emergent features is often used as evidence in debates over compositionality, since some interpretations of the compositionality principle prohibit such features.¹⁶ I am *not* concerned here with what these cases can tell us about compositionality. I will assume that some complex concepts are at least sometimes constructed in a way that grants them emergent features in order to see what this says about the status of concepts as a kind.

¹⁶ Compositionality is usually understood to demand that some property (or range of properties) of a complex concept is determined wholly by the properties of its constituents. For instance, the content of a complex concept is compositional iff it is determined entirely by the content of its constituents (plus their mode of combination, a qualification I will usually omit). Discussions of how concepts combine, and whether various theories of concepts are able to satisfy compositionality, has occupied much of the contemporary literature (Fodor & Lepore, 2002). Given the volume of ink that has been spilled on the topic, I will for reasons of space not take up the issue of compositionality here. Note though that the statement of the combination requirement says only that concepts should be *able to* combine to form larger structures in a productive and systematic way, not that they *always* do so. Attending to the modality of the compositionality requirement may provide an answer to many worries about whether particular theories of concepts can satisfy it; see Prinz (2002) and Robbins (2002) for development of suggestions along these lines.

Two commonly cited sources of information for emergent features are stored category exemplars and ‘background information’, including theoretical information about the categories being combined (Murphy, 2002, chap. 12; Prinz, 2002, chap. 11). In constructing the concept WOODEN SPOON we might draw on memories of experienced wooden spoons in order to derive the feature LARGE. Some intuitive, albeit fallible, psychological theorizing might lead to the conclusion that Harvard grads who become carpenters are unlikely to be materialistic. The point about these examples is that they involve processes that draw on multiple kinds of concepts, namely exemplars and causal models. If complex concepts are built up in part by drawing on these non-compositional sources of information, then conceptual combination will itself be a process that draws on multiple kinds of concepts, since both exemplars and theory-like causal models are among our concepts. So concept combination is another example of a process that can only be described in terms more general than those restricted to specific, narrow kinds of concepts.

This point can be illustrated by considering one particular model of concept combination, Costello and Keane’s (2000) Constraint-guided Conceptual Combination (C^3) model. In this model, nominal concepts are assumed to be frame-like structures consisting of features, which are described as attribute-value pairs. Combination involves constructing complex representations out of a set of input concepts plus a knowledge base of background information. The knowledge base, importantly, can consist of both stored exemplars and ‘theoretical’ information about the categories involved in the combination, and it is this information that is tapped to produce emergent features such as those mentioned above. Combinations are ranked for their diagnosticity, plausibility, and informativeness, and those combinations that meet a certain threshold of acceptability on these dimensions are retained.

The C^3 model illustrates two points that are important in the present context. First, since it assumes only that concepts are representations having frame-like structure, it can potentially combine *all* of the kinds of concepts canvassed here. Prototypes, exemplars, ideals, and causal models can all be (and usually are) understood as structures made up of features having attribute-value form. So the C^3 model (or an extension thereof) can potentially create complex concepts out of all of these different kinds of concepts. Second, the model illustrates how many kinds of concepts may contribute to conceptual combination by providing emergent features. Exemplars and theory-like information about a category are both drawn on in the current implementation of the model. Extending the knowledge base to include other information about the categories being combined should pose no problem in principle, since all of this information should be representable within frame structures. So the C^3 model illustrates a process for combining concepts independent of the particular representational subkind that they belong to, as well as a process that draws on a wide range of concepts.

3. *Common modes of acquisition.* Related to the previous observation is the fact that these different kinds of concepts may be acquired in ways that involve the same sorts of cognitive processes. Prototypes are frequently taken to be acquired by abstraction over experienced exemplars, but complex causal schemata may also be acquired through experience with exemplars. This process might involve causal reasoning or abductive reasoning. Language and other public representational media provide another route to concept acquisition. As Putnam (1970) observed, in telling someone the meaning of a term we often describe the prototype associated with it. This shows that prototypes can be acquired without abstraction over instances. The same is true of category ideals, which frequently lack instantiations and are culturally transmitted. Again, where there are generalizations subsuming multiple representational kinds,

we are justified in thinking that those kinds, although distinct in some respects, also belong to a higher-level kind.

Further, the way some concepts are acquired typically makes use of other kinds of concepts. On the standard prototype account, prototypes are acquired by abstraction from experienced exemplars. So acquiring a prototype depends on having previously acquired exemplars. On some Theory theorists' view, acquiring information to be used in the construction of causal models depends on having acquired statistically-organized information (Gentner & Namy, 1999; Keil, 1989). Normal acquisition of one kind of concept, on these accounts, depends on acquisition of another kind. This suggests that these separate kinds belong to a system containing processes and mechanisms that allow them to interact, not to completely disjoint subsystems.

One candidate for a process that draws on different kinds of representations in facilitating concept acquisition is structural alignment (Gentner, 1983; Gentner, 2003; Gentner & Namy, 1999; Gentner & Markman, 1997; Markman, 2001). Structural alignment is a process of comparing representations for their common features and the relations that link those features. Alignment enables comparison between representations that overlap only in certain respects. In formal models of structural alignment, local feature matches are computed first, followed by relations between local features, and finally higher-order relations. The output of this process can be used for various purposes, including to compute similarity judgments or make inferences from one domain to another. Once two representations have been aligned, candidate inferences are often spontaneously proposed across the categories being compared (Gentner & Markman, 1997). Structural alignment is, then, an example of a general process that operates over representations with imperfectly matched internal structures. All that it makes is the minimal

demand that, as in the case of the C^3 model described above, concepts be representable in a frame-like format containing features and relations between them.

4. *Long-term memory processes.* Concepts are stored in and recovered from long-term memory, and they are interconnected in ways that are geared towards permitting appropriate deployment and modification. The processes that govern encoding and retrieval of long-term memory information will accordingly need to manipulate all of the various kinds of concepts canvassed here, otherwise we will have no explanation for the fact that the system behaves in an adaptive and task-appropriate fashion.

Ideally, long-term memory should be organized in a consistent and minimally redundant fashion (Millikan, 1998, 2000). Consistency and non-redundancy are maintained through the subject's CAT concepts being linked in long-term memory. What *makes* CAT_1 , CAT_2 , etc. all cat concepts is that they refer to cats. What enables the thinker to *recognize* this fact is that they are connected to one another via identity links that indicate this co-reference relation to the thinker. The system needs some way of recognizing when concepts refer to the same things. It is not important whether these links are thought of on the model of semantic networks or some other formalism, as long as there is a mechanism that functions in the right way.

Such identity links might be established in processes of concept acquisition. Suppose that many concepts of concrete objects are acquired via encounters with instances of the categories that they denote. (Or public representations of category instances, as with pictures of tigers or toy elephants.) As these instances are encountered, exemplar representations are encoded and stored. As these exemplars are being stored, a prototype of the category as a whole begins to be generated. In virtue of being created by a process that takes a set of jointly stored exemplars as input, the prototype might be linked to them co-referentially. So prototype and exemplar identity

links could be formed as a result of the way that prototypes themselves are acquired. Prototypes, then, once acquired, could serve as the gateway to acquiring further concepts. When we receive information about the causal dependency structure of an object that we have categorized using a prototype, that information might be stored in long-term memory as a causal model of the category; and this relationship of co-reference is again stored via an identity link.

The mechanism at work in all of these cases is one of *referential chaining*. The operative rule might be: when a novel kind of information is acquired about something that falls under a concept C, and that information can be used for categorizing or reasoning about Cs in general, store that information as a concept C' that is co-referentially linked with C. One interesting possibility is that concepts acquired via the structural alignment mechanism described above might be linked co-referentially in memory. Structural alignment can promote inferences across aligned representations and encourage a search for properties that go beyond the ones so far observed. In this way, structural alignment could function not just in concept acquisition, but also as a mechanism for unifying information about categories acquired at different times and in different contexts.

While structural alignment is a well-attested process, its role in facilitating new ways to represent categories and establishing co-reference links is so far speculative. But whether or not this is the correct one, there must be some such mechanism at work to organize the contents of long-term memory. Processes that maintain a well-functioning memory generalize over different kinds of concepts: they treat these different kinds of representations as fungible to a certain degree. Concepts are stored, linked together, and retrieved by a set of memory processes that are indifferent to the subkind that they belong to. The existence of such a mechanism operating over

all kinds of concepts gives us reason to think that concepts themselves should be seen as a superordinate kind.

To sum up the moral of this section: there are both ‘top-level’ and ‘bottom-level’ demands at work when considering whether a category should be split. The bottom-level demands, such as more fine-grained predictions or taxonomies that more closely track the underlying mechanisms at work in a domain, tend to favor fragmenting a putative kind into subkinds. The top-level demands, such as generalizations that range over the subkinds, tend to favor unification. In the case of concepts, the existence of a set of common overarching processes and generalizations indicates that these subkinds are a more coherent and systematic object of study than their differences might otherwise lead us to think. The generalizations in the domains canvassed above are most perspicuously stated in terms of concepts *generally*, not particular subkinds of concepts.

6. Pluralism and hybrid models

The pluralist view of concepts presented here should be distinguished from hybrid theories of concepts. There are two ways of explaining what hybrids are. On one view, hybrids are theories on which there are multiple different kinds of conceptual representation, each of which is available to be activated for use depending on the nature of the case. Anderson and Betz (2001) implement such a model that incorporates both exemplar-based and rule-based categorization components in the ACT-R architecture. In their model, the system switches back and forth between using exemplars and using categorization rules derived from experience, depending on which is more successful in the task at hand. However, at no point does the system represent categories in a way other than via exemplars or via defining rules. No “mixed”

representations, incorporating aspects of each distinct process, are ever constructed. It is clear that, on this interpretation, hybrid theories are indistinguishable from pluralism, since they are similarly committed to the use of distinct kinds of conceptual representations in task-appropriate ways. I will set aside this sense of hybrids in the following discussion.

On another, more prominent view, hybrids are theories on which concepts are identified with *single* representations that possess two or more distinct components that have significantly different characteristics—i.e., they represent different information, are processed in distinctive ways, etc. These theories are at least *potentially* different from pluralism, since they posit single concepts, albeit large and highly structured ones, for each category the subject can entertain thoughts about. The earliest such models incorporated a “core”, which was often thought of as being like a classical definition, and a peripheral “identification procedure”, which was more or less prototype-like. Osherson and Smith (1981) suggest that the core contains defining properties and determines certain of a concept’s inferential relations to other concepts, while the identification procedure is employed in rapid categorization. This general strategy of dividing up the job of accounting for performance on particular cognitive tasks among distinct subcomponents of a single concept has been followed by most hybrid views.

One similar hybrid model is Smith, Shoben, and Rips’ featural model, which incorporates just such a set of definitional and stereotypical features (Rips, Shoben, & Smith, 1973; Smith, Shoben, & Rips, 1974). In the model, concepts (or, in their terms, the meanings of lexical items as encoded in semantic memory) are lists of features, some of which are “more defining” than others. Concepts are representations containing a definition-like part and a prototype-like part. In categorization and other semantic decision tasks (e.g., making typicality judgments), there are two stages of feature comparison. Suppose one is trying to decide whether penguins are birds. In

stage one, *all* features of PENGUIN, characteristic and defining, are accessed and compared to all features of BIRD. If the comparison meets a certain criterion value, then a categorization decision is returned. If not, then stage two of the process is activated. In stage two, only the defining features of the concepts are compared in order to render a judgment. This is a way of formalizing the intuitive idea that definitional cores represent a kind of fall-back procedure to be used when prototypical information fails. One finds similar tactics tried in other hybrids, although the particular division of labor is not always the same. In some theories it is fast vs. slow categorization that is at issue, or categorization judgments vs. typicality judgments. I will ignore such differences here.

As a general theory of concepts, the featural model inherits the same problems as the classical, definitionist theory and prototype theories do individually. The positive empirical evidence for definitions remains slim, and prototypes themselves are inadequate to fully explain categorization. Yet other models incorporate a theory-like component and a prototype component (Keil, 1994; Keil, Smith, Simons, & Levin, 1998) or an unstructured, atomistic component and a theoretical component (Rips, 1995), and obviously many other combinations are possible. But a key part of the argument for pluralism is that no single kind of representation so far considered is adequate to capture even most of the data. Consideration of the limits of these representations taken individually displays one source of pressure on hybrid models: relatively small hybrids are unlikely to account for much of the relevant range of empirical phenomena.

It would clearly be fruitless to argue against these hybrids one by one. Pluralism is related to hybrid models, in that it denies that any of the pure theories is monolithically correct. Yet these hybrid models still propose a “one size fits all” approach to concepts, because hybrids are

typically thought of as being employed in all circumstances when a concept is entertained and used. There is, then, a dilemma for hybrid accounts: either they fall short of accounting for some phenomenon or other, or they attempt to account for *everything* and grow to the point of being unmanageable.

There seems to be no way out of this dilemma for hybrids. A hybrid of two kinds of representations, neither of which is individually sufficient to cover all the data, will only have an advantage if the union of its constituent kinds covers all of the data. The natural response for a hybrid theorist would be to build into each concept *every* form of potentially relevant information. Concepts then might be thought of as single, structured bundles containing exemplars, prototypes, theoretical information, causal information, category ideals, etc. This representational inflation faces at least two problems. First, structures that include all of a person's ways of representing a category are too large to serve as units of thought. Working memory, where concepts are assembled on-line into thoughts, is has limited capacity, which in effect imposes architectural limits on the size of human concepts. Second, there is no reason to think that *all* category information is retrieved and used each time a person tokens a concept. The experimental data surveyed so far tend to show effects attributable to some relatively well-delimited subset of this information, rather than to the entire bundle. The question then becomes how to selectively retrieve *parts* of this large body of information to serve the immediate goals of the subject better. But again, this begins to collapse into the pluralist position. I would suggest completing the collapse and giving up the idea of finding a single kind of representational structure, even a *hybrid* kind, to serve all of our cognitive purposes.

7. Conclusion

I have argued that both the Singularity and Uniformity Assumptions are false, and hence that a theory of concepts should be a pluralist theory, in the sense that it should treat concepts as a set of different kinds of representational structures that are acquired and deployed under different circumstances and for different ends. Nevertheless, this does not prohibit us from formulating generalizations over concepts at higher levels of abstraction. Examples of these generalizations include those that involve the logical form of complex thoughts, conceptual combination, common mechanisms of concept deployment and concept acquisition, and the organization of long-term memory. The subkinds of concepts canvassed here belong to a mosaic of such generalizations that criss-cross one another, some subsuming the entire set of subkinds, some only a few of them. Insofar as there are such generalizations, concepts retain their unity as a higher-level cognitive kind. Hence pluralism should be viewed as a theory *of* concepts (or at least the framework for such a theory), not a spur to eliminate concepts.

The focus of research within a pluralist framework should be on spelling out the precise ways in which context can promote the use of one or another of these kinds of concepts. Normal cognitive development undoubtedly plays a role; children's concepts may differ from those of adults in systematic ways (Opfer & Siegler, 2004; Poling & Evans, 2002). The precise task at hand also clearly has an influence. Whether we are judging typicality, categorizing a perceived object, or reasoning counterfactually influences how we represent a category (Rips, 1989). Ontological domains may also exert a broad influence on the kinds of concepts we construct to represent them. Within the object domain, for instance, many have held that artifacts and natural kinds are represented differently (Estes, 2003); outside of the domain of objects, little is yet known about how concepts of events, persons, and abstract objects might be structured. These all remain open avenues of investigation.

The core argument for pluralism rests on the functionalist considerations sketched at the outset. It gains support from the general inadequacy of the competing monolithic views. There is at this point little hope that a single monolithic theory of concepts will have the scope to account for all of the data. But there is nevertheless compelling evidence that each of the kinds of structure posited captures *some* subset of the data. Many kinds of cognitive structures play the concept role. These representations do not belong to autonomous, disjoint systems. Rather, they constitute different aspects of the human conceptual system.

Acknowledgements

First, my gratitude to two anonymous referees for this journal for their thorough, incisive, and extremely thoughtful comments. Their efforts improved the paper significantly. Thanks also to Bill Bechtel, Mason Cash, Edouard Machery, Gualtiero Piccinini, Jesse Prinz, Philip Robbins, Dan Ryder, Steven Sloman, Eric Winsberg, and audience members at the Society for Philosophy and Psychology (Winston-Salem, NC, 2005) and the Southern Society for Philosophy and Psychology (Durham, NC, 2005) for helpful comments on earlier versions of this paper and for discussion of many of the points developed here.

References

- Ahn, W., & Dennis, M. J. (2001). Dissociation between categorization and similarity judgment: Differential effect of causal status on feature weights. In U. Hahn & M. Ramscar (Eds.), *Similarity and categorization* (pp. 87-107). Oxford: Oxford University Press.
- Ahn, W., & Kim, N. S. (2001). The causal status effect in categorization: An overview. *The Psychology of Learning and Motivation*, 40, 23-63.

- Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, 41, 361-416.
- Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin and Review*, 8, 629-647.
- Bailenson, J. N., Shum, M. S., Atran, S., Medin, D. L., & Coley, J. D. (2002). A bird's eye view: Biological categorization and reasoning within and across cultures. *Cognition*, 84, 1-53.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, 11, 211-227.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 629-649.
- Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Block, N. (1997). Anti-reductionism slaps back. *Philosophical Perspectives*, 11, 107-133.
- Boster, J. S., & Johnson, J. C. (1989). Form or function: A comparison of expert and novice judgments of similarity among fish. *American Anthropologist*, 91, 866-889.
- Boyd, R. (1991). Realism, anti-foundationalism, and the enthusiasm for natural kinds. *Philosophical Studies*, 61, 127-148.
- Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and*

- intellectual factors in categorization* (pp. 141-174). Cambridge: Cambridge University Press.
- Burnett, R. C., Medin, D. L., Ross, N. O., & Blok, S. V. (2005). Ideal is typical. *Canadian Journal of Experimental Psychology*, 59, 3-10.
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24, 299-349.
- Craver, C. F. (2004). Dissociable realization and kind splitting. *Philosophy of Science*, 71, 960-971.
- Dehaene, S. (1997). *The number sense*. Oxford: Oxford University Press.
- Estes, W. K. (1994). *Classification and cognition*. Oxford: Oxford University Press.
- Estes, Z. (2003). Domain differences in the structure of artifactual and natural categories. *Memory and Cognition*, 31, 199-214.
- Fodor, J. A. (1994). Concepts: A potboiler. *Cognition*, 50, 95-113.
- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Oxford University Press.
- Fodor, J. A., & Lepore, E. (2002). *The compositionality papers*. Oxford: Oxford University Press.
- Fodor, J. A. & Pylyshyn, Z. (1988). Connectionism and Cognitive Architecture. *Cognition*, 28, 3-71.
- Gelman, S. A. (2003). *The essential child*. Oxford: Oxford University Press.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.

- Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195-236). Cambridge: MIT Press.
- Gentner, D., & Namy L. L. (1999). Comparison in the development of categories. *Cognitive Development*, 14, 487-513.
- Gentner, D., & Markman, A. B. (1997). Structure-mapping in analogy and similarity. *American Psychologist*, 52, 45-56.
- Gopnik, A., Sobel, D. M., Glymour, C., Schulz, L. E., Kushnir, T, & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 3-32.
- Hacking, I. (1991). A tradition of natural kinds. *Philosophical Studies*, 61, 109-126.
- Hampton, J. A. (1995). Similarity-based categorization: The development of prototype theory. *Psychologica Belgica*, 35, 103-125.
- Heit, E., & Barsalou, L. W. (1996). The instantiation principle in natural categories. *Memory*, 4, 413-451.
- Hirschfeld, L. A., & Gelman, S. A. (1994). *Mapping the mind: Domain specificity in cognition and culture*. Cambridge: Cambridge University Press.
- Johnson, K. E. (2001). Impact of varying levels of expertise on decisions of category typicality. *Memory and Cognition*, 29, 1036-1050.
- Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 924-941.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge: MIT Press.

- Keil, F. C., Smith, W. C., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition*, 65, 103-135.
- Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research*, 52, 1-26.
- Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin*, 112, 500-526.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social categories: The role of causal reasoning. *Cognitive Science*, 14, 551-577.
- Lakoff, G. (1987). Cognitive models and prototype theory. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 63-100). Cambridge: Cambridge University Press.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 3-81). Cambridge: MIT Press.
- Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory and Cognition*, 28, 41-50.
- Machery, E. (2005). Concepts are not a natural kind. *Philosophy of Science*, 72, 444-465.
- Machery, E. (2006). How to split concepts: A reply to Piccinini and Scott. *Philosophy of Science*, 73, 410-418.
- Malt, B. C. (1989). An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 539-555.

- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40, 230-262.
- Markman, A. B. (2001). Structural alignment, similarity, and the internal structure of category representations. In U. Hahn & M. Ramscar (Eds.), *Similarity and categorization* (pp. 109-130). Oxford: Oxford University Press.
- Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, 129, 592-613.
- Marr, D. 1982. *Vision*. New York: W. H. Freeman.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London*, 200, 269-294.
- Medin, D. L., & Atran, S. (1999). *Folkbiology*. Cambridge: MIT Press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
- Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? *Annual Review of Psychology*, 51, 121-147.
- Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, 32, 49-96.
- Millikan, R. G. (1998). A common structure for concepts of individuals, stuffs, and basic kinds: More mama, more milk, and more mouse. *Behavioral and Brain Sciences*, 22, 55-65.
- Millikan, R. G. (1999). Historical kinds and the "special sciences". *Philosophical Studies*, 95, 45-65.

- Millikan, R. G. (2000). *On clear and confused ideas: An essay about substance concepts*. Cambridge: Cambridge University Press.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge: MIT Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700-708.
- Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 3-27.
- Opfer, J. E., & Siegler, R. S. (2004). Revisiting preschoolers' *living thing* concept: A microgenetic analysis of conceptual change in basic biology. *Cognitive Psychology*, 49, 301-332.
- Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9, 35-58.
- Palmeri, T. J., & Gauthier, I. (2004). Visual object understanding. *Nature Reviews Neuroscience*, 5, 291-304.
- Piccinini, G., & Scott, S. (2006). Splitting concepts. *Philosophy of Science*, 73, 390-409.
- Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, 343, 263-266.

- Poling, D. A., & Evans, E. M. (2002). Why do birds of a feather flock together? *British Journal of Developmental Psychology*, 20, 89-112.
- Prinz, J. J. (2002). *Furnishing the mind*. Cambridge: MIT Press.
- Putnam, H. (1970). Is semantics possible? In H. Putnam, *Mind, language, and reality: Philosophical papers, Vol. 2* (pp. 139-152). Cambridge: Cambridge University Press.
- Ratneshwar, S., Barsalou, L. W., Pechmann, C., & Moore, M. (2001). Goal-derived categories: The role of personal and situational goals in category representation. *Journal of Consumer Psychology*, 10, 147-157.
- Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1141-1159.
- Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, 50, 264-314.
- Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). Cambridge: Cambridge University Press.
- Rips, L. J. (1995). The current status of research on concept combination. *Mind and Language*, 10, 72-104.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1-20.
- Robbins, P. (2002). How to blunt the sword of compositionality. *Nous*, 36, 313-334.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.

- Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22, 189-228.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge: Harvard University Press.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81, 214-241.
- Smith, J. D., & Minda, J. P. (2003). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 3-27.
- Tarr, M. J. (2003). Visual Object Recognition: Can a Single Mechanism Suffice? In M. A. Peterson and G. Rhodes (eds.), *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes* (pp. 177-211). Oxford, UK: Oxford University Press.
- Tarr, M. J., & Bühlhoff, H. H. (1998). Image-based object recognition in man, monkey, and machine. *Cognition*, 67, 1-20.
- Weiskopf, D. A. (forthcoming a). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research*.
- Weiskopf, D. A. (forthcoming b). First thoughts. *Philosophical Psychology*.
- Weiskopf, D. A. (2007). Concept empiricism and the vehicles of thought. *Journal of Consciousness Studies*, 14, 156-183.
- Whittlesea, B. W. A., Brooks, L. R., Westcott, C. (1994). After the learning is over: Factors controlling the selective application of general and particular knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 259-274.

Yamauchi, T., & Markman, A. B. (1998). Category learning by inference and classification.
Journal of Memory and Language, 39, 124-148.