PHILOSOPHICAL PSYCHOLOGY, VOL. 17, NO. 3, SEPTEMBER 2004

Routledge
Taylor & Francis Group

Concepts á la modal: an extended review of Prinz's *Furnishing the mind*

ARTHUR B. MARKMAN & C. HUNT STILWELL

ABSTRACT *In* Furnishing the mind, *Prinz defends a view of concept representation that assumes all representations are rooted in perception. This view is attractive, because it makes clear how concepts could be learned from experience in the world. In this paper, we discuss three limitations of the view espoused by Prinz. First, the central proposal requires more detail in order to support the claim that all representations are modal. Second, it is not clear that a theory of concepts must make a realist assumption. Third, the arguments focus on object categories that can be described by features, which are only one of many types of categories. Despite the flaws in the book, however, it clearly highlights a road that can be taken by those interested in defending an empiricist view of concepts.*

1. Preliminary points

A central question in cognitive science concerns the nature of concepts. Cognitive agents have the ability to recognize, classify, communicate about, reason about, and interact with a variety of objects, individuals, events, and abstract entities in the world. By having concepts, people are able to recognize cases in which a particular instance of a concept ought to be treated as a member of a more general class, and are thus able to bring their previous experience to bear on new situations.

Prinz's book, *Furnishing the mind*, aims to provide a philosophical account of concepts that is based on recent proposals from psychology about the role of perceptual information in cognitive representations (e.g. Barsalou, 1999). There is a real appeal to invoking perceptual representation in our understanding of concepts. Our ability to perceive the world is marked by great flexibility to adapt to changes in the conditions of the world. Our ability to form mental images (including visual, auditory and even tactile images) also seems to have this flexibility. Furthermore, our ability to use concepts seems to have a similar degree of flexibility, which allows us to apply our past experience to many new situations. Thus, perceptual representations seem to have just the sort of characteristics that a theory of concepts requires.

In psychology, the jury is still out about whether a theory of concepts based solely in perceptual representations will provide a sufficient explanation of concep-

Arthur B. Markman, Department of Psychology, University of Texas, 1 University Station, A8000, Austin, TX 78712, USA, email: markman@psy.utexas.edu

ISSN 0951-5089/print/ISSN 1465-394X/online/04/030391–11 © 2004 Taylor & Francis Ltd
DOI: 10.1080/0951508042000286749

TABLE 1. Desiderata used by Prinz to evaluate theories of concepts

Scope	A theory of concepts must accommodate the range of concepts people possess. Concepts are real entities that share one or more properties.
Intentional content	Concept representations stand in for extramental things.
Cognitive content	Concepts are individuated both by reference to extramental things and also by their connection to other mental representations.
Acquisition	Theories of concepts must explain how concepts can be acquired. They must also be compatible with some reasonable proposal about how humans come to have concepts.
Categorization	Theories of concepts must explain how items in the world are identified as members of a particular concept.
Compositionality	Theories of concepts must explain how two or more concepts may be combined to form more complex concepts.
Publicity	Theories of concepts must explain how two or more people may come to share the same concept.

tual processing. As a philosophical argument, however, we find the perceptual account of concepts offered by Prinz to have significant shortcomings. In this extended review, we first summarize the key points of Prinz's account. Then, we discuss three aspects of Prinz's proposal that we felt were problematic: (1) the basis of concepts, (2) the specific support for the claim that concept representations are tied to perceptual modalities, and (3) and the relationship between psychological data and philosophical theory.

2. Prinz's perceptual proxytype proposal

Prinz begins the book with a list of desiderata against which he evaluates philosophical theories of concepts. The list of desiderata is shown in Table 1. We will have more to say about certain of these desiderata later, but for now, we should point out that versions of these criteria have been used in many philosophical arguments about concepts.

In the three chapters that follow, Prinz discusses the limitations of a number of previous proposals for concepts. These chapters are meant largely as a review of previous work for people new to theories of concepts. Chapter 2 discusses classic theories of concepts based on images and definitions. Chapter 3 examines prototype and exemplar views that were popular in psychology starting in the 1970s. Finally, Chapter 4 discusses the "theory theory"—a view that suggests that concepts are defined by theories that explain why items are members of particular categories—and Fodor's (1998) atomic view of concepts. Prinz argues that none of these alternatives effectively addresses all of the desiderata, and thus, the rest of the book presents his own attempt at a theory that can do so.

There are two principles at the core of Prinz's theory of concepts. The first is *concept empiricism*, which he defines as: "All (human) concepts are copies or combinations of copies of perceptual representations" (p. 108). Because perceptual representations are constructed as people interact with the world, perceptual

representations reflect experience with the world and thus are not innate. The second is *modal-specificity*, which assumes that "Concepts are couched in representational codes that are specific to our perceptual systems" (p. 119). Thus, Prinz denies the existence of concept representations that are not specificially attached to human perceptual systems.

Finally, a critical component of Prinz's theory is the notion of a *proxytype*. A proxytype is a specific perceptual instance of a concept that is generated when reasoning about concepts. The perceptual instance may involve more than one modality (e.g. the proxytype for a dog in the context of thinking about pets might involve both a visual image of the appearance of the dog as well as an auditory image of the sound of the dog's bark, and perhaps even a tactile image of the feel of the dog's tongue as it licks your hand). Different proxytypes may be generated in different contexts, so that the dog proxytype will be different when thinking about pets than when thinking about *The call of the wild*. Furthermore, people have a flexible ability to create mental images, and this flexibility can (presumably) be used to generate images of concepts whose instances have never been encountered directly. Barsalou (1999) refers to this process as *simulation*.

The rest of the book is devoted to fleshing out the notion of a proxytype relative to the desiderata presented in Chapter 1, and defending this view against a variety of objections. To this end, Prinz marshals both psychological evidence and philosophical argument. As an example of the former, Prinz points to the studies of the simulation process to support the use of perceptual representations in concepts. For example, Wu and Barsalou (in preparation) find that people asked to generate properties of a lawn rarely say "roots," but people asked to generate properties of a rolled-up lawn often say "roots." This finding suggests that perceptual availability increases the likelihood that a property will be listed.

On the philosophical side, Prinz discusses how proxytypes can refer to natural kinds when they contain perceptual information. For example, he suggests that the proxytype for the element GOLD might contain perceptual information that it is shiny, yellow, and malleable. These properties also refer to the perceptually represented ideas shiny, yellow, and malleable, which (presumably) are in the head. Yet, the concept itself is assumed to refer to the natural kind (GOLD) in the world. He argues that the properties are used to allow a person to track a member of a natural kind, and that the real kind in the world is used to ensure that the internal representation remains in correspondence with the world. As he says, "Conceptual development is a matter of fine-tuning our abilities to track real contents via appearances."

Obviously, we cannot do justice to Prinz's entire argument in a few pages. Nonetheless, we wanted to give a flavor of the elements of the book. In the following sections, we discuss three ways in which we feel the arguments fall short of providing a basis for a theory of concepts. These arguments will address the foundations (i.e. the desiderata) of the discussion of concepts, the psychological basis of the theory, and the effectiveness of the philosophical arguments.

3. Do we have access to what (if anything) is out there?

Table 1 lists the desiderata that Prinz specified for evaluating theories of concepts.

These criteria are presented as crucial for the success of any theory of concepts, because they define what it is that such a theory is supposed to do. In our view, this set of desiderata makes two problematic assumptions. First, this view follows much of the philosophical literature in assuming a property-based view of concepts. Second, this view makes a strong realist assumption that is unwarranted. We discuss each of these issues in turn.

The standard view of concepts in cognitive science is that there is some set of properties that determines whether an item is a member of a category. Theories disagree about issues such as whether the properties are theories, essences (a set of necessary properties, e.g. a particular genetic code for animal kinds), or characteristic properties (in which exemplars are related to the category based on their similarity to the stored category representation). Nevertheless, all of them seem to agree that concepts are described by properties (although, for Fodor, 1998, the set of properties that describes a concept is empty).

This view of concepts effectively characterizes many of the kinds of concepts that typically form examples in research papers in cognitive science. For example, GOLD seems to be effectively described by some set of characteristic properties (that includes being shiny, yellow, and malleable) and/or some set of essential properties (a particular chemical makeup). That said, there are many other kinds of concepts that we use on a regular basis that do not fit this mold. For example, consider the concept GAME. From Wittgenstein (1953) forward, it has been clear that there are no necessary and sufficient sets of properties that determine whether something is a game. It is possible to argue that there is simply a prototype for GAME and that new exemplars are evaluated relative to that prototype based on their similarity to it. However, a deeper analysis suggests that GAME is actually defined by the role it plays relative to other concepts, namely, as the second argument to the relation play (x,y) (Markman & Stilwell, 2001); concepts like GAME defined by their position within a relational system are called *role-governed categories*. Other examples of role-governed categories are JOB, BARRIER, and BEVERAGE (Markman & Stilwell, 2001; see McRae *et al.*, 1997, for further discussion of role-governed categories).

Other concepts seem to denote the relation among other items. For example, kin terms (e.g. SISTER, UNCLE, FATHER-IN-LAW) are relational concepts of this type. These relational categories are also poorly characterized as referring to sets of properties. The upshot of this discussion is that assuming property-based concepts leads to one of two potential problems. One possibility is that a theory of concepts will attempt to account for role-governed and relational categories in the same way as property-based categories, in which case the theories are likely to falter on these kinds of categories. A second possibility is that other types of concepts such as role-governed and relational concepts will be defined away as some other kind of entity, in which case the theory of concepts evaluated by these desiderata will be too narrow.

The perceptual view of concepts is likely to have particular difficulty with relational and role-governed categories. The relations and roles that define many categories are not specifically perceptual (nor do category members share many perceptual properties). Prinz (see also Barsalou, 1999) makes an interesting move to deal with this type of structure. Mental representations require specifying both the representational medium as well as a set of processes that act on the representation to use the information in it (Markman, 1999; Palmer, 1978). The non-perceptual information is shunted to the processes that operate on the representation. In particular, the way the perceptual approach deals with many arguments about the importance of amodal representations is to assume that they can be dealt with by the simulation process that combines representations.

There are two potential problems with this move. First, the simulation process has never been specified in enough detail to evaluate whether it would in fact solve the problems for which it has been invoked. Perhaps more importantly, however, it is not clear how the simulation process can know what it is supposed to do without some representation that guides it. Fodor (1998) makes a related point by discussing the computational problems that emerge for theories of concepts because some decision procedure must be defined that determines what processes use the information in a concept representation.

A second critical assumption in the desiderata is that concepts refer to actual extramental entities. This realist assumption is adopted consciously, and Prinz spends a few pages defending it in the first chapter, as well as periodically throughout the rest of the book. Despite the importance of realism in many philosophical theories of concepts and meaning, this assumption seems superfluous and unempirical, and it introduces a number of additional problems to be dealt with that could be avoided without it.

Rather than making a realist assumption, it would be easier to adopt a *coherence-based* framework (Dietrich & Markman, 2000). That is, the only information that any person has about the outside world comes from perceptual representations, which are themselves mental entities. Thus, rather than being concerned with whether a particular concept correctly refers to all and only proper extramental entities, it would be better to generate a theory in which the use of the concept attempts to remain consistent with other representations in the system. Such a system would actually implement the spirit of Prinz's proposal for the GOLD concept discussed above, without jumping through philosophical hoops to ensure that the concept adequately refers to an extramental entity. A coherence-based framework would also avoid problems engendered by twin-Earth thought experiments (e.g. Putnam, 1988). A being on a planet exactly like ours in which water has the chemical structure XYZ rather than H_20 would still have the same concept of water as we do prior to the dissemination of scientific discoveries about chemical compositions, because the same concept would create coherence among that being's representations as it would on this planet regardless of what is "actually" true in the real world.

The impetus for the realism requirement seems to be twofold for Prinz. First, he argues that for people to communicate about a concept (the publicity desider-

atum), they must have identical concepts. Prinz assumes that the only two ways that publicity can be satisfied are assuming that concepts are innate or that they reflect real entities in the world (so that everyone acquires the same set of categories). Because Prinz rejects nativism in Chapter 8, realism seems to be the only option left open for him.

However, there is a third route to publicity, one that Prinz touches on occasionally without realizing its relevance to publicity. On the empiricist view, our concepts are acquired through interactions with the world. These interactions take place in a social context, and concepts are largely delimited not by the structure of the "real" world, but through the pattern of references that people make to concepts in conversation (Markman & Makin, 1998). Take as an example the concept "CONCEPT." Cognitive scientists with a wide range of theoretical beliefs about the nature of concepts, and thus a wide variety of concept representations, are able to communicate more or less effectively about concepts because we all acquired these concepts of "CONCEPT" within the same communicative context.

The second impetus behind the realist assumption for Prinz is the belief that in order for perceptually represented concepts to correctly refer to objects with hidden (i.e. non-perceivable) essences, the world must be structured such that perceptual properties are highly correlated with perceivable essences. Because the essentialist assumption is itself a realist assumption, it thus suffers from the same problems, and thus can also be avoided if we recognize that concepts are acquired in a social context. This social context allows individuals to benefit from the perceptual experience of a large group of individuals, and thus map concepts onto correlations between perceptual properties that may not be readily observed through the limited experience of a single individual, and may or may not map onto something essential about the real-world referents of those concepts.

In summary, two core desiderata start Prinz's theoretical program off on a problematic footing. First, he assumes that all concepts are property-based. Second, he assumes a strong realist stance. Prinz is in good company in making these assumptions. Nonetheless, they limit the likelihood of success of his theoretical enterprise.

4. Can a case be made for modal representations?

The core of this book is the idea that mental representations are best characterized as being tied directly to perceptual modalities. This proposal permits Prinz to adopt an empiricist account of concepts. While the perceptual apparatus clearly has innate structure, the specific items that people perceive are the basis of experiential knowledge. Thus, if concepts are rooted in perceptual information, they must necessarily be rooted in experience rather than in some innate structure.

There are a number of objections to the idea that all concepts have a perceptual basis. For example, there are many concepts that an individual never directly experiences. Some of these are fictitious items (e.g. unicorns), while others fall outside of a typical person's experience (e.g. sloths, electrons, and craters on the dark side of the moon). In addition, there are concepts that do not seem to have

obvious perceptual components (e.g. virtue, truth, and justice). After describing his theory of concepts, Prinz addresses many of these objections.

In order to defend the perceptual view of concepts, Prinz must give a detailed description of his proposal and then provide a convincing account that perceptual information can indeed handle the objections that have been raised in the past to image-based proposals for concepts. Unfortunately, Prinz's account falls short on both of these requirements. The modal approach is not sufficiently detailed to support firm arguments that mental representations are modal. Furthermore, some of Prinz's arguments appear to contradict the perceptual basis of concept representations. In this section, we first address two representative places where the argument is not sufficiently detailed. We then turn to places that appear to run counter to the spirit of the modal approach.

4.1. The need for more detail

Prinz argues that perceptual representations are compositional on the basis of links among different perceptual representations (either within the same modality or across modalities). As one example of a link, Prinz suggests that hierarchical perceptual structure can be represented using a link that recommends that one representation be replaced with a second when one "zooms in on" (p. 145) the representation. As a second example, when the same object is observed in more than one configuration (e.g. dog sitting and panting or running across a field), then these representations can be linked together by a transformational link, "Because each is stored as a permissible transformation of the other" (p. 145).

Prinz simply assumes that these links can be implemented as perceptual representations. However, the system has to have some sort of control structure that determines what sort of link should be followed. Furthermore, the transformations among representations must have some description that states what the transformation is. That is, when we form a mental image of a dog sitting, we know we are looking at a dog sitting. When we then imagine the same dog running, we know the dog is running. It is not clear how this is done with just perceptual information. That is not to say that it cannot be done, but rather that the details have been skipped.

These details are important for evaluating this proposal against the desiderata. The links in the previous paragraph are meant to ensure that concepts are compositional. For example, Prinz suggests that knowing that playing fetch will make a dog happy along with knowledge that happy dogs wag their tails is sufficient to form an image of a dog playing fetch and wagging its tail. While there is no doubt that perceptual information is used to generate a mental image of a dog wagging its tail, it is not at all clear how this compositionality is carried out using only perceptual information. The proposal cannot be evaluated because it is not presented in sufficient detail.

As a second example, Prinz argues that infants may begin to develop a notion of causality based on their sensory-motor interactions with the world. At some level, this suggestion is likely to be correct. For example, there is ample evidence that infants gain substantial knowledge about their physical and social world only after

they begin to move through the world (for a review, see Campos *et al.*, 2000). This knowledge arises as a result of the experience of self-directed movement rather than simply maturation, as non-locomotor infants who are given a simple locomotor device that they can control begin to show conceptual changes that infants who are not given this opportunity for self-exploration do not achieve.

Nonetheless, the fact that the concept of causality is rooted in perceptual experience does not require that the concept have only a perceptual component. In particular, it is not clear what the perceptual content of the actual notion of causality would be. It is not enough to have the appearance of a causal event. The prototypical example of billiard balls colliding involves contact between the balls, and an immediate effect of one ball on another. Still, we can have a sequence of images of objects that interact without feeling it necessary to invoke a causal attribution. Furthermore, we talk about many situations as being causal that do not have this kind of simple perceptual content. For example, we can say that the assassination of the Archduke Ferdinand caused World War I, even though there is no apparent physical or spatial contiguity that are involved in this causal event. Indeed, the concept of causality is bound up with people's ability to form explanations, which seems to go beyond the sort of information that is likely to be represented perceptually (Keil & Wilson, 2000).

Prinz attempts to deal with difficult cases like this by suggesting that we can ground our abstract concepts in specific perceptual information. For example, we might represent the concept of virtue by "simulating acts of altruism and charity (e.g. running in front of a truck to save a child, handing a check to a person in need)" (p. 178). As appealing as this idea is, because abstract concepts are strongly contextually governed, it is not clear how one could simulate a virtuous act without first having a concept of virtue. Furthermore, we should note that this is another example in which a difficult case for the perceptual account is solved by assuming that the processes that operate on the representations will take care of the problem.

These two examples suggest that the details of the perceptual account must be presented before a philosophical account of concepts can be evaluated against the desiderata. Furthermore, the explanation of some of the hard cases for Prinz's view cannot be evaluated without specific information about how the theory accounts for these hard cases. As in so many things, God (or perhaps the Devil) is in the details.

4.2. Contradictory cases

Prinz defends the perceptual view of concepts against a number of objections. Some of the arguments are quite successful. For example, Prinz does a nice job of demonstrating that a perceptual account of concepts provides a good basis for empiricism. Unfortunately, there are a number of arguments in the book that appear to defend the perceptual view by adopting an amodal view of representation and then simply assuming that these amodal representations are actually perceptual

As one example, Prinz invokes a version of Putnam's (1988) division of linguistic labor to explain how people have concepts of entities that they have not experienced directly. Presumably, non-experts know that one concept is related to

others, but the actual perceptual content that grounds the concept is only possessed by a select few. One can easily do a thought experiment from this point, in which everyone shares a concept that they believe is grounded in someone else's perceptual representation, despite the absence of anyone who has such a perceptual grounding. Thus, the existence of an expert seems strictly irrelevant for actually possessing the concept.

In order to evaluate whether this distribution of perceptual information is problematic it is necessary to examine the basis of the semantics of mental representations. On the realist view that Prinz adopts, the key aspect of concept representations is that they permit people's concepts to stay in correspondence with the extramental entities they represent. As long as experts have perceptual representations that ensure that concepts appropriately refer to the right extramental entities, the absence of detailed perceptual information in most people is not problematic.

However, if most individuals possess a concept without having the perceptual basis for it, then there must be some other kind of information in the representations that permits the concepts to be used appropriately. Typically, this function of concepts is carried out by assuming that concepts are also defined in part by their connections to other concepts. (That is, concepts have both a correspondence and a functional role component to their semantics; see Dietrich & Markman, 2003.) Of course, the functional role view of concepts assumes that many (if not most) concepts are defined by their connections to other concepts.

On the surface, this web of conceptual connections seems to be amodal. Prinz argues that these concepts are verbal, and thus have the perceptual information associated with the words that generate them (e.g. sounds and images of words). On the one hand, this account does save the empiricism that is central to Prinz's view of concepts. On the other hand, the existence of a large body of concepts that are perceptual only in their relationship to words seems to stretch the perceptual account thin. Furthermore, there are a number of good reasons not to identify concepts with words. For one, languages have many synonyms that seem to map onto the same concept (e.g. doctor and physician). In addition, we have many polysemic words that can be used to refer to many similar but non-identical concepts. For example, the word *newspaper* can refer to the type of paper on which newspapers are printed, to a physical daily paper, to the organization that publishes the paper, or to the content of the newspaper (e.g. Klein & Murphy, 2001). If concepts are not simply represented by words, then it is not clear how concepts defined by functional role are based in perception.

Finally, when Prinz distinguishes between modal and amodal representations, he points out that there is much evidence from neuroscience that multiple perceptual modalities converge on some areas of the brain, and so there are cells that are not sensitive to any particular perceptual input. He suggests that these areas of the brain may serve as multi-modal convergence zones, which may later be used to reactivate those perceptual areas. This description of the activity of cells that are not tied to particular modalities seems like an exercise in where to draw the line between modal and amodal representations rather than a principled distinction between perceptual and non-perceptual representations. After all, the brain has many areas

that are devoted to single perceptual areas. Furthermore, there is a high degree of connectivity among parts of the brain. Thus, it would be difficult to find any area of the brain that was not connected at all to some area whose primary function was some kind of perceptual processing. To suggest that all such areas still involve perceptual representations seems simply to rename amodal representations "multimodal representations."

This argument suggests that the distinction between modal and amodal representations may be less important to concepts than Prinz (and others such as Barsalou) make it out to be. Empiricism is an important element of a theory of concepts, but it is not clear that empiricism requires that all aspects of concepts be tied to perception. Human experience is rooted in perception, but the ability to connect across modalities (and ultimately to use symbols that are so far removed from perception as to make their connection to perception tenuous at best) is crucial for developing concepts. Thus, clinging rigidly to the idea that all concept representations are perceptual seems to cause more problems than it solves.

5. So, is it safe to mix philosophy and psychology?

The desiderata in Table 1 clearly lay out a philosophical research program. The core criteria describe issues that have occupied the philosophical literature on concepts for hundreds of years. A book evaluating psychological theories of concepts would have a very different set of criteria for evaluating the success of a theory (e.g. Murphy, 2002). In psychology, desiderata would emphasize both consistency with existing data and ability to generate ideas for new experiments.

The pragmatic use of psychological theories enables them to tolerate a certain degree of vagueness, because even a theory that is missing key details can be used to generate ideas for new studies. Furthermore, a gap in a theory itself may lead to new experiments. For example, Barsalou's (1999) description of the perceptual symbol systems approach to concepts appeals to perceptual simulation processes that generate concept representations on the basis of contextual factors. These simulation processes are not defined. Nonetheless, the basic structure of the perceptual symbol systems approach has been sufficient to motivate a number of new studies.

Philosophical arguments are less tolerant of vagueness about the details. This account of concepts is weakest at the point where the details have not been presented to a degree that allows the theory to be evaluated. So, what is the philosophy equivalent of the pragmatic utility that theories play in psychology? We suggest that this book makes the world of analytic philosophy safe for empiricism again. Until recently, empiricism has been limited to areas of philosophy that are largely ignored by analytic philosophers (e.g. Bergson, 1988; Merleau-Ponty, 1964). However, over the last decade, there has been a resurgence of empiricism across cognitive science. For example, Elman *et al.* (1996) provided a spirited defense of empiricism using connectionist models as a vehicle. Despite its flaws, if *Furnishing the mind* plays a similar role within the philosophy community, then it will have served a valuable function.

References

BARSALOU, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–660.

BERGSON, H. (1988/1911). *Matter and memory*, N.M. PAUL & W.S. PALMER (Trans.). New York: Zone Books.

CAMPOS, J.J., ANDERSON, D.I., BARBU-ROTH, M.A., HUBBARD, E.M., HERTENSTEIN, M.J. & WITHERINGTON, D. (2000). Travel broadens the mind. *Infancy*, *1*, 149–219.

DIETRICH, E. & MARKMAN, A.B. (2000). Cognitive dynamics: computation and representation regained. In E. DIETRICH & A.B. MARKMAN (Eds) *Cognitive dynamics* (pp. 5–29). Mahwah, NJ: Lawrence Erlbaum.

DIETRICH, E. & MARKMAN, A.B. (2003). Discrete thoughts: why cognition must use discrete representations. *Mind and Language*, *18*, 95–119.

ELMAN, J.L., BATES, E.A., JOHNSON, M.H., KARMILOFF-SMITH, A., PARISI, D. & PLUNKETT, K. (1996). *Rethinking innateness*. Cambridge, MA: MIT Press.

FODOR, J.A. (1998). *Concepts: where cognitive science went wrong*. New York: Oxford University Press.

KEIL, F.C. & WILSON, R.A. (Eds) (2000). Explaining explanation. In *Explanation and cognition* (pp. 1–18). Cambridge, MA: MIT Press.

KLEIN, D.E. & MURPHY, G.L. (2001). The representation of polysemous words. *Journal of Memory and Language*, *45*, 259–282.

MARKMAN, A.B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum.

MARKMAN, A.B. & MAKIN, V.S. (1998). Referential communication and category acquisition. *Journal of Experimental Psychology: General*, *127*, 331–354.

MARKMAN, A.B. & STILWELL, C.H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, *13*, 329–358.

MCRAE, K., FERRETTI, T.R. & AMYOTE, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes*, *12*, 137–176.

MERLEAU-PONTY, M. (1964). *The primacy of perception, and other essays on phenomenological psychology, the philosophy of art, history, and politics*. Evanston, IL: Northwestern University Press.

MURPHY, G.L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

PALMER, S.E. (1978). Fundamental aspects of cognitive representation. In E. ROSCH & B.B. LLOYD (Eds) *Cognition and categorization* (pp. 259–302). Hillsdale, NJ: Lawrence Erlbaum.

PRINZ, J. (2002). *Furnishing the mind: concepts and their perceptual basis*. Cambridge, MA: MIT Press.

PUTNAM, H. (1988). *Representation and reality*. Cambridge, MA: MIT Press.

WITTGENSTEIN, L. (1953). *Philosophical investigations*, G.E.M. ANSCOMBE (Trans.). New York: Macmillan.

WU, L.L. & BARSALOU, L.W. (in preparation). Grounding concepts in perceptual simulations: evidence from property generation.