Embodied Cognition and the Percept/Concept Distinction

Michael Hayward
Department of Cognitive Science, UCSD

December, 1998

Disclaimer: This paper is the result of a brief research foray undertaken as part of a degree program requirement. Although I believe it to be entirely accurate, it is neither complete nor authoritative. Please do not cite.

Embodied Cognition and the Percept/Concept Distinction

Michael Hayward
Department of Cognitive Science, UCSD
mhayward@cogsci.ucsd.edu

Embodied cognition, a theoretical framework which extends the definition of cognitive activity to include bodily interactions with the world, stands as a challenge to the conventional view of cognition as primarily abstract, rule-based symbol manipulation in the brain. Proponents of embodied cognition downplay or altogether deny the existence of such an abstract conceptual system independent from perception. This paper integrates theory and evidence across many domains in cognitive science, in an attempt to evaluate the claim of some embodied cognition proponents that modal, perceptual activity can account for most or all of the cognitive phenomena traditionally viewed as "conceptual". I find that much of the evidence presented as support for either side is misleading, as it is based on the assumption that existing theories are far more complete than they really are. Nevertheless, I conclude that there is enough strong evidence to support several key claims, including the existence of perceptual components in conceptual thought, and the inability of some classes of propositional conceptual system to adequately explain the full range of human conceptual phenomena.

1. Introduction

1.1) Two Perspectives on Cognition

In the last few years, cognitive science has seen a resurgence of interest in placing cognition back in the context of a body acting in a social and material environment. In *Being There*, Andy Clark stresses that "the biological mind is, first and foremost, an organ for controlling the biological body" (pg.1). Beyond simply emphasizing the importance of body and environment, proponents of this view argue that cognitive behavior *cannot* be understood without recognizing that our everyday cognitive activity extends outside of the brain, and into our bodily interactions with the world (Hutchins, 1995). Inherent in this "embodied cognition" view is an appreciation for the importance of the sensorimotor system as a locus of cognitive activity, rather than just a mediator between the world and the mind.

The embodied cognition approach stands in contrast to many of the assumptions that have underpinned cognitive science for decades. The process view supported by these (typically implicit) assumptions is as follows: humans receive perceptual input, model it in some abstract "language" of conceptual thought, apply rule-based symbol manipulation to process it, and finally produce motor output as a result. With the digital computer as a metaphor, problem solving and planning are offered as the quintessential cognitive activities under this "disembodied" view. Understanding the representations which underlie cognition then reduces to finding the appropriate conceptual language to support just this kind of symbol manipulation.

In contrast, embodied cognition advocates seek to describe cognition through representations formed dynamically across the interactions of the mind, the body, and the world. Researchers downplay or, in extreme cases, entirely deny the existence of an abstract conceptual symbol system independent of perception. Instead, they posit perceptual symbol structures, often instantiated in the same brain regions responsible for representing external stimuli. The challenge is therefore to redescribe conceptual phenomena in perceptual terms, and it is this "percept/concept debate" which is the focus of this paper. Just how far can a purely perceptual mind go towards explaining cognition?

To give a sense for the appeal of each approach, let me briefly review some of the intuitive arguments which support them. On the one hand, it seems self-evident that the brain *did* evolve primarily to control bodily behavior, and thus sensorimotor activity should form the foundation of the mind. Furthermore, we certainly experience the *conscious* feeling of manipulating mental images, as if “pictures” truly were a medium of higher thought. On the other hand, it's not clear how abstract concepts like "justice" could possibly be represented in a modal, perceptual form. Furthermore, classical, propositional artificial intelligence has produced some powerful techniques which, in select problem domains, fit human performance data very well. These are the kinds of issues which have lead to the debate over concept representation.

It is misleading, however, to suggest that there exists a single dimension along which these views differ. By posing the problem as a "debate", it may seem to suggest that there are only two opposing views. In fact, there is an enormous range of theoretical positions presented in the literature. In order to make this discussion tractable, I have tried to reduce the key dimension to that of the role of the sensorimotor system in conceptual processing. Inevitably, however, such an approach fails to entirely capture the depth and significance of all the issues involved.

1.2) Approach

The concept literature is filled with unsubstantiated assertions and speculation. Debates are dominated by arguments over definitions and partitionings of the theoretical space, rather than empirical evidence (e.g., the "situated cognition" special issue of Cognitive Science centered around Vera & Simon, 1993).

My goal is to try to provide a somewhat more objective, evidence-oriented analysis of this issue, integrated across many domains in cognitive science. I will do so by testing the explanatory power of the major frameworks against a common set of conceptual phenomena, each of which has been offered as evidence for some view in the literature. I thus explore each theory's ability to explain a wide range of phenomena, including those that are most and least natural for it.

While attempting an objective evaluation of the evidence, the balance of my arguments must nevertheless focus on confirming or denying the claims of the embodied camp. This imbalance is necessary primarily because the embodied approach is challenging the established view. As a result, most of the debate and discussion is focussed on supporting or refuting that challenge. Rarely is the disembodied approach argued for directly; instead, it is supported by rejecting the claims of the challengers.

1.3) Outline of the Paper

- In Section 2, "Differing views of the role of perception...", I revisit the different sides of the debate, giving a brief review of the background and important dimensions along which the research varies.

- In Section 3, "A Functional View of the Human Conceptual System", I define "concept" by considering the types of phenomena theorists are trying to explain. The phenomena selected have each been presented as evidence for some view in the literature. This is the list of phenomena against which the major frameworks will later be judged.

- In Section 4, "Domains and Theories", I outline the wide range of methods and domains of research that are relevant to the issue. In each, I try to summarize the contributions of a few influential individuals whose work I have chosen to focus on.

- In Section 5, "Analysis", I present each of the phenomena from Section 3, and consider how well they are explained by each of the major frameworks from Section 4. In the course of this analysis, I present empirical evidence, and explore issues raised in section 4 in greater depth.

- In Section 6, "Conclusions", I summarize the results of the analyses and present my own conclusions with regard to the debate, and the state of the research as a whole.

2. Differing views of the role of perception, the body, and the environment

In this section, I give a more detailed characterization of the views held on each side of the debate. Though the extreme versions of these positions differ greatly, none of the views described can be dismissed outright by the available evidence. Influential researchers in modern cognitive science represent all sides.

2.1) The Disembodied, Conceptual View

The disembodied view is an implicit assumption underlying much of modern cognitive science. The advent of the digital computer had a tremendous impact on our view of the human brain. The field of artificial intelligence was not just passively influenced by psychological research; it actively influenced psychology through its emphasis on serial processing, input/output models, theory of general computation through symbol manipulation, etc. Computers became very powerful, very quickly through these techniques. With the adoption of the view of "cognition as computation", cognitive science looked to these same frameworks for guidance.

In computers, it was clear that the "real work" was being done in the CPU; the "perceptual systems" were simply modified versions of old technology: TV sets, typewriters, cassette players, etc. Artificial intelligence adopted the same view, and set about trying to model the "real" intelligent behaviors of human beings: problem solving, theorem proving, chess playing, and so on. The goal was to find the appropriate internal representation to model and manipulate the necessary symbols; input and output were secondary.

Newell and Simon's influential "physical symbol system hypothesis" (1976) made the disembodied ideal explicit: intelligence could be completely described in terms of the correct symbols moving in the correct ways, independent of their implementation. With this view, the body and the world are pushed into the background. AI met with important successes with this approach, developing processes (e.g., production systems) and representations (e.g. slot-and-filler structures, semantic nets, etc.) that are still among the most powerful computing techniques applied today. It was in this environment that cognitive science looked for propositional representations to model human knowledge, and rule-based symbol manipulators to model reasoning. Many of these approaches were founded on the theory of formal logics, and the mathematics of computers became the mathematics of the mind.

Today, after decades of such work, many of these same views still dominate artificial intelligence (e.g., Lenat & Guha (1990)) and cognitive science (e.g. Fodor & Pylyshyn (1988)). I will be referring to this approach as "disembodied", although the term is rarely used explicitly.

The nature of human cognition under this view, then, is inherently nonperceptual. Concepts are typically represented as propositional structures abstracted from the sensory modalities by some process of filtering and transduction. In this way, a computer-like serial processing sequence is enforced: perceptual input is received, translated for conceptual processing, and then translated into motor signals for output. The conceptual representations form a sort of "language" underlying all cognitive activity. In their most traditional form, these representations explicitly include dictionary-like necessary and sufficient conditions (e.g. the concept bachelor is explicitly defined as "unmarried male") or invariant feature lists (e.g. bachelor: gender(male) + status(unmarried)). Other modern forms include probabilistic components and fuzzy set theory to account for graded category membership, and other such techniques. The common tie between all of them is that concepts are treated as abstract (non-perceptual), well-defined entities manipulated through a set of explicit rules.

2.2) The Embodied, Perceptual View

The embodied approach emphasizes a completely different aspect of cognitive activity altogether. In this view, the social and material environment are not only the driving forces behind cognition, but actually form an integral part of most cognitive activity. Perceptual and motor activity are in no way secondary to more abstract processes; perception and action in the real world form the foundation of cognition. Cognitive activity is defined by a constant, dynamic relationship between representations within the brain and outside the brain. In many cases, imagistic, analogistic, and dynamical explanations replace logical symbol manipulation as the foundation for human reasoning. Importantly, representations often take the most "external" form possible, in order to save the time and energy required to translate information into "deeper", more abstract forms for processing. When possible, this means that representations are left outside of the body entirely, in the form of "cognitive artifacts", as described by Hutchins (1995). Kirsh & Maglio (1992) present a clear example of pushing representation out into the world in their research on the video game "Tetris". According to Kirsh & Maglio, when it becomes necessary to rotate a simple geometric shape as a basis for a strategy judgement under strict time constraints, expert players prefer physical, rather than mental, rotation of the figures. This seems to suggest that the sensorimotor loop is in some way more efficient than "deep" conceptual processing.

In this view, the next best thing to external representations is to manipulate internal perceptual traces directly, without having to translate to a deeper "conceptual" language. Conceptual symbols are: (a) systematically, rather than arbitrarily, related to the perceptual traces they represent, and (b) active in the same modality-specific perceptual region of the brain as the original perceptual trace. I will explore this view in depth in Section 4.

2.2.1. Why Embodiment?

Given the significant successes of the disembodied approach, what caused researchers to look to an embodied view as an alternative? Numerous reasons were behind the move for change. Experimental psychologists were discovering that even the most robust human behaviors are highly context-dependent; relatively minor changes to otherwise highly controlled laboratory experiments could change results significantly. Even phenomena as "low-level" as color perception were subject to context effects (Thompson, 1995). As a result, the validity of the laboratory setting came into question, as researchers formed a new appreciation for the importance of the complicated and noisy environments in which humans operate naturally.

Along with a renewed appreciation for the importance of context, it soon became clear that organisms, and particularly humans, create and take advantage of structure in their environments to offload internal cognitive work (Hutchins, 1995). Furthermore, these environmental interactions are so quintessentially cognitive that it is hard to justify arbitrarily limiting the term "cognitive" to refer to internal activity.

In AI, it was becoming clear that tasks such as theorem proving, originally assumed to be the most difficult, were turning out to be easier than the most basic physical behaviors like object recognition, voice recognition, and dynamic balance. For some, this suggested that these tasks are "effortless" for humans precisely because they were the first, and most important, abilities for early humans. To model human intelligence, one must take into account what the system was originally designed to do. Furthermore, "toy" problems were being solved efficiently using classic AI representations, but these systems were rigid and brittle: in the face of noise and unpredictable changes (hallmarks of the real world), the systems would fail catastrophically.

Finally, some philosophers exploring the nature of human meaning were asking difficult questions about models of understanding based entirely on symbol manipulation. In short, if abstract symbols were given meaning only through relation to other abstract symbols, where did those *other* abstract symbols get their meaning? To avoid infinite regress, some argued, meaning must ultimately be grounded in the only "concrete" processes available: interaction with the world.

2.2.2. Additional Themes in Embodied Cognition Theory

Along with an increased appreciation for the role of the environmental context, the body, and the perceptual system, embodied theories regularly share some other features. None of these features is strictly necessary for, nor limited to, an embodied view, but tend nevertheless to be common to much of the work.

Dynamics. With an emphasis on real-world interaction, real-time behavior becomes an integral part of cognitive theory. More than simply acting "as fast as possible", there are issues of synchrony between internal and external events, as well as physical limitations of signal propagation times to take into account (see Port & van Gelder, 1995, for a review). In many cases, researchers abandon classical representations in favor of a lower-level, dynamical systems view of the mind and its interaction with the world. Note that this does not necessarily entail anti-representationalism (Port and van Gelder remind us that "dynamical systems can be representational without having their rules of evolution defined over representations", p.12), nor does it necessarily preclude the possibility of a symbol-manipulation interpretation of the same mechanism at some higher level of abstraction.

Emergence. Emergence is difficult to define precisely, but with the renewed emphasis on representations distributed across the traditional brain/body/world boundaries, finding the locus of cognitive phenomena becomes more difficult. Rather than trying to identify the hierarchically control mechanism responsible for cognitive activity, the embodied view focuses on identifying dynamic systems of interaction from which phenomena "emerge". Rodney Brooks' "subsumption architecture", a robotic design methodology in which behavior emerges from the interaction of semi-autonomous sensory systems, is a classic example I will review in Section 4.

3. A Functional View of the Human Conceptual System

In this section, I pursue the definition of "concept" in greater depth. Finding the available definitions highly variable and contentious, I move on to a functional view of the types of conceptual phenomena we are trying to explain. The resulting list of phenomena will later be used to judge the explanatory power of the various theories (in Section 5).

3.1) Defining "concepts"

A brief survey of the philosophical literature on conception reveals that there is no accepted definition at all; the available options are vague, highly variable, and contentious. To give an impression of the range, the following list is a small sample of definitions taken from various dictionaries of philosophy:

- "psychologically or logically what we grasp in understanding an expression"
- "any abstract notion or idea by virtue of which we apply general terms to things"
- "a semantically evaluable, redeployable constituent of thought, invoked to explain properties of intentional phenomena such as productivity and systematicity"
- "that aspect of mentality which, ideally, refers to an object, property or relation and specifies some properties of that item"
- "an accumulation of related facts and/or processes to which a symbol (representational or abstract) is assigned"

Any finer-grained definitions than these are even more controversial, and too tightly intertwined with particular theories. It is abundantly clear that there is no definition precise enough for rigorous scientific analysis, and any attempt on my part to construct such a definition would be misguided. The only reasonable way to proceed, then, is to ignore the intentional definitions altogether, and *functionally* identify which phenomena we are aiming to describe. In the next section, I have compiled some of the decisive phenomena that have guided individual researchers in their exploration of the issue.

3.2) Conceptual Phenomena

The following list presents some of the human cognitive phenomena that are relevant to our debate. I have chosen to concentrate on those phenomena which have been presented in the literature as the strongest examples of supporting evidence for either side of the debate. I will then judge all of the frameworks against the entire collection of phenomena, seeing each at its (purportedly) best and worst. I use this approach to achieve some distance from the carefully constructed world of supporting evidence that each author uses to present his or her work.

As a result of the selection process, the following list does not capture a complete, cohesive view of the conceptual system (if such a thing even exists). Instead, it is a disparate sampling, biased somewhat towards the "fringe" of what might be considered conceptual. Nevertheless, any complete theory of the human perceptual/conceptual system would need to be able to explain each and every one.

- Mental simulation of spatial events

(searching a room in your "mind's eye", mentally rotating an object, etc.)

- Using concepts that are heavily context-dependent

(e.g., "normal", "fair", or "large", as in "large molecule" vs. "small planet")

- Reasoning with ad hoc categories

(e.g., items in the room that are likely to get stolen)

- Understanding, reasoning with, and extending words without clear definitions

(e.g., making sense of the sentences "Greg opened the store", but not "Greg opened the tree". Also, extending understanding of the phrase "open the coconut" to understand "open the walnut")

- Reasoning with abstract concepts

(e.g., comparing the concepts of "justice" and "fairness")

- Reasoning about impossible events

(e.g., reasoning that dragons weigh less on the moon than on earth)

- Integrating conceptual and physical activity

(e.g., finding the most "expensive", or "dangerous", item in a room)

- Rule-based categorization

(e.g., using a rule to categorize 323,456,391 as an odd number, despite never having seen it.)

- Subtle perceptual categorization

(e.g., distinguishing male and female faces, without necessarily knowing how we do it.)

Also, we need to account for:

- Phylogenetic development: how we developed smoothly from, presumably, purely perceptual ("stimulus-response") creatures

- Ontogenetic development: how (and whether) we develop from concept-free or concept-limited infants to highly conceptual adults

While it's true that any complete theory of conceptual activity would need to account for each item on this diverse list, virtually no work in the field today claims to be such a complete theory. Nevertheless, it is worthwhile to bring compare the diverse frameworks in an effort to determine which are at least *consistent* with the phenomena.

4. Domains and Theories

In this section, I explore the relevant domains of research, outlining the key issues, individuals, and theories in each. Needless to say, the domain boundaries are fuzzy; there is a lot of overlap. Due to space limitations, I have had to omit potentially relevant domains of research and individuals. I have chosen one or two influential researchers from each domain, based on their importance and on satisfying the goal of surveying as wide a range of viewpoints as possible in a limited amount of space. For the most part, I will limit my review to a description of the work, without introducing evidence (which I will present in Section 5, "Analysis").

I begin with Table 1, which summarizes the information I'm about to present: the major domains, key researchers, and key terms I'll explore. Table 1 will hopefully clarify the scope of the analysis, and will provide a reference point for the later sections of the paper.

Table 1: Key Researchers, Domains, and Theories

Researcher(s)	**Relevant Interest/Theory**	**Notes**
Barsalou	Perceptual Symbol Systems	“Simulators”. Modal, schematic, dynamic
Fodor	Abstract Conceptual Systems	Propositional, computational, atomistic concepts
Johnson	Image-Schemas	Amodal, perceptual, metaphorical, dynamic
Kosslyn	Mental Imagery	Picture-like, part of a dual system
Lakoff	Theory-based Categorization	Concepts defined in terms of naïve theories
Rosch	Categorization	Basic categories, objective categories
Neisser	Categorical Perception	Culturally-biased “direct perception”
Harnad	Symbol Grounding Problem	Abstract system, but meaning grounded in perception
Brooks	Mobots/Subsumption Architecture	Embodied mobots, anti-representational
Beer	Dynamical Systems	Real-time action, internal states allowed
Piaget	Ontogenetic Origins of Concepts	Separate conceptual system at age 1.5
Mandler	Ontogenetic Origins of Concepts	Separate conceptual system from birth
Clark & Grush	“Emulators” as Origins of Symbols	Phylogenetic origins

4.1) Perceptual Symbol Systems

This subsection is rather unique, in that it deals specifically with the work of one particular researcher: Larry Barsalou. I felt that it was appropriate to dedicate a section exclusively to Barsalou's theory because it is one of the most current, complete, and influential views in this domain, and because it directly addresses many of the major issues in this paper. The remaining subsections identify more general domains of research, spanning multiple people and theories.

Barsalou's (in press) "Perceptual Symbol System" framework is perhaps the most comprehensive description of a purely perceptual, modal implementation of a conceptual system. In his view, conceptual processing involves filtered, schematic perceptual symbols which are manipulated in the same perceptual regions of the brain dedicated to processing external stimuli. Specifically, a perceptual symbol is a subset of the unconscious portion of neural activity associated with a perceptual state. Thus, all such symbols are "modal", in the sense that they are tied to the specific sensory modality in which they were originally active. Barsalou makes the strong claim that these modal symbols alone account for all conceptual activity; that the brain does not support any form of amodal conceptual content. Examples of the kinds of content captured by perceptual symbols include: "tiger", "stripedness", "above", "comparison", "anger", etc. He stresses that these modal symbols are *not* like conscious mental images, but are abstracted perceptual *schemata*. Through a process of selective attention that breaks down and abstracts perceptual traces into symbols, the schemata can come to represent abstractions like shapes independent of orientation. Further details on the abstraction mechanism are not provided.

In his description of human conceptual processing, Barsalou invokes the concept of a "simulation": a dynamic reactivation and recombination of perceptual symbols in working memory. He stresses the active, productive nature of the perceptual system, in contrast to the "passive recording device" view. Because the schemata are componential and feature-like, they can be recombined in an infinite number of ways, accounting for the effectively infinite human conceptual capacity.

In order to include action and affect as elements with which to ground meaning, Barsalou includes proprioception and introspection as "perceptual systems". Introspection provides the "perceptual" medium necessary for generating concepts referring to cognitive operations (e.g. "comparison") and emotional states.

Barsalou's system is much like a perceptual re-cast of a classical symbol-manipulation approach. The real difference, though, is that the symbols remain *modal*; instead of transducing perceptual input into amodal forms for processing in other parts of the brain, Barsalou contends that traces merely need to be "filtered down" into a schematic form, and processed entirely within the sensorimotor system.

4.2) Philosophy

Not surprisingly, the oldest tradition in concept research comes from the domain of philosophy. Much work is still at least partly philosophically-oriented, perhaps because the paucity of clear scientific evidence leaves much to conjecture. In this section, however, I try focus only on "classically" philosophical arguments. Even so, it would be easy to get bogged down in the regress of thousands of years of philosophical debate; such a review could easily fill volumes. My goal here is simply to touch on two directly relevant issues: the philosophy of propositional concepts as represented by Jerry Fodor, and the philosophy of meaning, as represented by the work of Steven Harnad and Mark Johnson.

4.2.1) The Philosophy of Concepts as Abstract Symbol Systems

I take this opportunity to briefly revisit the views of advocates of disembodied concepts, who, for reasons discussed earlier, appear underrepresented in this paper as a whole. I use Jerry Fodor's (1998) views as a model because he represents one of the most outspoken and influential advocates of the propositional viewpoint.

In brief, Fodor argues adamantly for a view of cognitive activity as computation. The style of computation he advocates is classic, Turing-style symbol manipulation. The goal of his work is an "atomistic" theory of concepts: one in which concepts are distinct, discrete entities with precise modes of interaction defined. Importantly, Fodor stresses that the structure of complex mental representations determines the kinds of possible interaction allowed. Promoting his view that "mental representation is a lot like language" (p. 25), he places perhaps more emphasis on the syntactic structures than on the atomic concepts that fill them. Propositional-style structured representations carry most of the interesting content in this view; concepts themselves are simply arbitrary atoms to be bound to the available structure slots.

Fodor is well known for seeking what he terms a 'language of thought' with combinatorial syntax and semantics. The propositional content (i.e. meaning) of a sentence in this language corresponds to the function that determines, for any possible situation, whether that sentence is true or false. As for the perceptual nature of concepts, he explains that his framework, and any like it, "demand mental representations that aren't images" (p. 9).

4.2.2) The Philosophy of Meaning

Ultimately, any complete theory of human conceptual ability is necessarily also a theory of meaning; after all, concepts function primarily as units of understanding. While the philosophy of meaning is an enormous topic, I am choosing to concentrate on one specific, particularly relevant, aspect known as the symbol grounding problem. The symbol grounding problem is relevant to our debate because it represents an attempt to bridge the gap between conceptual *structure* and *function*. The central issue is how mental representations become endowed with *meaning*.

To see the significance of the problem, I turn to Steven Harnad's (1990) analysis. Harnad describes a classical, propositional style system of abstract symbol manipulation, and asks: how is the *meaning* of arbitrary symbols to be grounded in something other than just more meaningless symbols? Harnad describes the problem as tantamount to trying to learn to speak Chinese using only a Chinese/Chinese dictionary: a "merry-go-round" of replacing symbols with symbols, none of which hold any more meaning than the previous ones.

Harnad's solution is to retain a very classical view of conceptual symbol systems (consisting primarily of symbol strings), but to ground the symbolic system in a non-symbolic one. The non-symbolic system is composed of "iconic representations", which are sensory projections of objects and events, and "categorical representations", which are feature detectors that pick out the "invariant features of objects and events" from their sensory projections. The argument is one that is held in various forms by most researchers in the field: that ultimately, only real objects, events, and actions in the world can hold true *meaning* for humans. By proposing a mediating system between the world and the abstract symbol system, Harnad preserves many of the basic features of the classic, disembodied view, while providing a mechanism for grounding meaning in such a system.

For Mark Johnson, it is unnecessary to work to preserve the classical abstract symbol system. Johnson agrees that meaning must be grounded in perception and action, but argues that the conceptual system should fundamentally operate in this meaningful domain, rather than having meaning attached through associations to a distinct system. Thus, Johnson dismisses the abstract "symbol strings" view and replaces it with a system of "image schemas", amodal but still fundamentally percept-like in nature. As Johnson's work is primarily language-oriented, I detail his influential views in the next subsection.

4.3) Language

Language has played an important role in concept research, because it is in language use that our concept definitions show up most explicitly. The relationship between language and the conceptual system is hotly debated. On the one hand, we have seen that Fodor argues that mental representations are in many ways isomorphic to language, and tightly bound to it through shared structural relations. On the other end of the spectrum, Barsalou confers no special status on language. He treats spoken and written language just like any other perceptual traces; to be abstracted and stored along with other perceptual symbols.

Mark Johnson (1987) falls between these two extremes, stressing that linguistic meaning is just one type of meaning. He cautions that looking at meaning through sentences may mislead us to consider only propositional-style theories. Nevertheless, he looks to language, or more specifically the *use* of language, as the primary source of evidence regarding the nature of the human conceptual system. Johnson points to pervasive use of metaphor in language, and suggests that this phenomenon is an indication of the fundamentally analogistic nature of concepts. This is an issue I will return to in the Analysis section, when I consider evidence.

For Johnson, there is more to understanding than just semantics: understanding is "historically and culturally embedded, humanly embodied, and imaginatively structured" (p. 175). Concepts are not objective; they derive their meaning only in the context of experience. Words, in turn, only get their meaning through their *use* in describing the world.

Johnson argues that propositional thought is only possible by virtue of a network of nonpropositional schematic structures derived from bodily experience. These structures, which he terms "image schemas", are not unlike Barsalou's perceptual symbols. They are "dynamic patterns" abstracted from perceptual traces, representing concepts like "containment" and "force". Image schemas are not as rich as an image, nor as abstract as a proposition. Unlike Barsalou, however, Johnson's schemas can be multimodal or amodal without being "nonperceptual". In his words, they represent "a recurrent pattern, shape, or regularity in ongoing activities" (p. 48).

4.4) Mental Imagery

Alongside the perceptual schemas of Barsalou and Johnson, it is necessary also to briefly explore the considerable literature on "mental imagery". With this term, I'm referring to the conscious experience of fairly complete, percept-like (e.g. "picture-like") traces in the absence of corresponding external stimuli. There is considerable debate about the significance of the role of such phenomena in conceptual processing. Perhaps the name most closely associated with mental imagery in the last two decades is that of Steven Kosslyn (e.g., Kosslyn, 1976). Kosslyn is a strong advocate of "offline" perceptual activity, but he does not claim that the processing of "pictures" can provide a foundation for all of human conceptual activity. Instead, he supports a dual-system view, in which an imagistic component interacts with a more traditional propositional system. The difficulty, of course, comes in having to explain precisely how these systems interact: how they communicate, how they "divide up" the cognitive workload, etc. This kind of dual-code system is not an uncommon view in the field, but there is, as yet, no clear consensus on the issue of interaction. I explore Kosslyn's mental imagery evidence in greater depth in the Analysis section.

4.5) Categorization

Categorization research is as close as experimental psychology comes to the direct study of the nature of concepts. There is no consensus on the precise nature of the relationship between the two, however. Categories have been presented, in one view or another, as: a special case of concepts, equivalent to concepts, and overlapping, but not subsumed by, concepts. Nevertheless, framing concept work in terms of categories allows researchers to pursue behavioral experiments based on category extensions: similarity judgements, category membership judgements, etc.

In one "realist" view, the world is divided into categories based on objective features. Rosch's (1978) foundational work on categorization took a realist approach, defining an objective set of category levels across all humans. The most salient level of abstraction for a category, for instance, is known as the "basic level", in which members of a category look alike and involve similar physical interactions. Such a system would have important consequences for concept research, suggesting that there is an objectively correct way to structure human knowledge. While this view embraces some of the major themes of embodiment (perceptual emphasis, and opportunistic use of structure in the world), it departs from the notion of context-dependence, and a considerable body of evidence has thrown the existence of objective category levels into doubt.

George Lakoff (1987) suggests that such a static view of categories, and concepts, is doomed to fail. Categories change according to context, and category membership is far from binary. As an example of graded membership, consider the seemingly well defined category "bachelor", taken to mean "unmarried male". Although it seems like any person will clearly fall inside or outside of this category, it's nevertheless the case that most people view James Bond as a "better" example of a bachelor than the Pope. Lakoff's solution to dynamic category definition is to embed all concepts in *theories*. Systematic relationships among concepts are explained through the internal structure and coherence of people's "naive theories" of how the world works. The naive theories, of course, are ultimately derived from experience. Thus, Lakoff argues, category membership judgements come not only from attribute similarities, but through reasoning about the causal relations among features.

Finally, Neisser (1987) takes a more heavily perceptual approach, invoking the idea of Gibson's direct perception in the context of a theory of categorical perception. An entire literature on "categorical perception" exists, concentrating most heavily on very low level perceptual categories (like colors) (see Harnad, 1987, for review), and mostly independent of research on the kinds of high-level conceptual categorization which are the focus of this paper. Neisser (and Clark, 1997, and Johnson, 1987) argues that simple "direct perception" is not sufficient to explain categorization, but that perception must be taken in the context of culturally-driven cognitive models. Neisser does not commit himself to a strict, perceptual view of concepts like Barsalou's, but he does make an important point about the evolutionary genesis of the conceptual system: even if we started as perceptual creatures and "worked in" towards an increasingly abstract, and ultimately distinct, conceptual system, we would still expect to be able to learn about the constraints on the conceptual system by studying the perceptual system. This brings us to the next important branch of the literature: development.

4.6) Development

Any complete theory of concepts would be considerably strengthened by an accompanying description of how the system develops. There are two timescales to conceptual development: the phylogenetic (evolutionary) scale and the ontogenetic (personal developmental) scale. Once again, the discussion of such matters could easily fill volumes, but I will provide only the briefest review of the key issues.

Jean Piaget, one of the most influential figures in human ontology, adopted a classical, propositional view of the adult conceptual system. He did not, however, believe in the innateness of concepts, instead treating infants at birth as "purely perceptual creatures". His task, therefore, was to explain the "leap" from perception to conception. Through extensive behavioral studies of infants, Piaget (1952) concluded that there was no activity before (roughly) the age of 1.5 years that could not be accomplished directly in the perceptual system. After that age, however, infants performed recall tasks which suggested they were able to represent objects and events in their absence, and thus had accomplished a conceptual level of thought. Piaget thus concluded that the distinct conceptual system appeared around age 1.5. He suggested that the origin of the conceptual system was related to the internalization of sensorimotor schemas, but was unable to describe the mechanisms more precisely.

More recently, Mandler (1992, 1997) revisits Piaget's views on the matter, suggesting that they have been largely unquestioned in the intervening decades. Based on considerable experimental evidence since Piaget's time, Mandler concludes that it is difficult to explain why conceptual development is so long delayed. Although she retains the assumption that the perceptual and conceptual systems are indeed distinct, she proposes that the two systems in fact develop in parallel from birth. I will revisit Mandler's evidence in the analysis portion of this paper.

Recently, Clark & Grush (in press) have attempted a further exploration of the kind of "internalized sensorimotor schema" Piaget described, in the context of trying to understand the phylogenetic origins of concepts. In principle, however, much of their description might just as easily apply to ontogeny. Based on evidence we will consider later, Clark & Grush provide a compelling description of a developmental process which begins with pure sensorimotor emulation of external events. Through small, but plausible stages, the emulation process gradually becomes increasingly detached from the external stimuli. In the first stage, the emulation simply mirrors the activity in the world. Later, the emulation takes place in the presence of the stimuli, but may run faster internally and thus "predict" the outcome of the external events. In this stage, there is regular interaction between the internal and the external. In later stages, the emulation can occur in the absence of external stimulation, a situation which Clark & Grush identify as the first "true representation". While their explanation ends there, I will return later to the intriguing possibility that such a mechanism might blend nicely into one of the perceptual schematic views we will analyze.

4.7) Robotics and Computer Simulation

Finally, I consider robotics and computer simulation to play an important role in understanding the nature of conceptual representation, because they allow us to implement and test otherwise heavily speculative theories in a concrete way. One of the obvious difficulties in studying concepts in human activity is lack of access; brain imaging techniques are not sophisticated enough to give us low-level descriptions of neural representation. In computers, however, representation is accessible, if not necessarily easily interpretable.

An interesting question arises for embodiment proponents when computer simulation is introduced as a tool. If real worlds, and real bodies, are essential components of human cognition, what level of detail (if any) is appropriate for modeling? If our perceptual systems are the locus of much cognitive activity, would research with robots be teaching us about general cognition, or about the specifics of the sensory apparatus we happen to endow them with? I can't tackle these issues in significant depth, given space restrictions. Nevertheless, it seems relatively uncontroversial to conclude that a proponent of disembodied symbol manipulation would likely be more optimistic about the possibility of capturing human-like cognitive activity in computers than a diehard supporter of the embodied view. This is not to suggest, however, that embodied theorists must completely dismiss the role of computer simulation, simply that they must necessarily accept additional limitations.

Rodney Brooks of the MIT Artificial Intelligence Lab has long been one of the most outspoken advocates of a "situated" view of cognitive activity within the robotics community. In his view, computers "have poisoned our view of what intelligence must be" (1991, p.). As an example of computation-biased design, Brooks points to what he refers to as the "sense-model-plan-act", or SMPA, design that dominates robotics. These systems employ the classic input/output view described above as one of the identifying features of the disembodied approach. Brooks takes a dim view of these approaches that "over-model"; in his words, "the world is its own best model" (1991,p. 583). Brooks' own robotic designs de-emphasize representation, instead opting to produce motor commands directly from perceptual input (e.g. working with 2-D visual input, rather than trying to model the 3-D world first). When representation is required, Brooks prefers "deictic representations": instead of modeling the world "objectively", his robots model entities exclusively in terms of their relationship to them (a "self-centered" view).

Brooks' design methodology, referred to as the "subsumption architecture", reflects most of the major themes of embodied cognition: situatedness, intelligence through interaction with the world, emergence, dynamics, etc. Brooks' robots are conglomerations of semi-autonomous sensorimotor subsystems, arranged loosely in hierarchical layers. Layers can encourage, interrupt, or override one another, but cannot "communicate" in any formal way. The complex, intelligent behaviors of the robot emerge from the interactions of sensorimotor subsystems, without requiring a central ("conceptual") modeling component.

Randall Beer (1995) also explores the theory behind the construction of situated agents, but focuses more heavily on seeking explanations in dynamical systems theory. Beer stresses that the primary goal for an agent is "appropriate behavior at the appropriate time". If real-time behavior is an essential component of embodied action, Beer reasons, dynamical systems theory is the most suitable tool for mathematical description. Beer's work is largely consistent with that of Brooks, though he is not as vehemently anti-representational. He argues that purely reactive systems (ones with no internal state) have no true autonomy; clearly internal states *do* exist in humans and other autonomous creatures. The question for Beer is how to dynamically couple those internal states to external states, and how to best make opportunistic use of structure available in the environment.

5. Analysis

In this section, I review the list of conceptual phenomena presented in Section 3. For each, I examine evidence from the domains described in the previous section, and consider the success of the major frameworks in explaining the phenomena. Once again, for reasons already discussed, I will focus most of my efforts on examining the perceptually-oriented frameworks, to see if they stand up to the strong claims made about them. As we'll see, in many cases the evidence (for either side) is not nearly as decisive as it is often made out to be.

Each of the following analyses exposes one or more relevant issues. Recall, the phenomena were chosen primarily because they have each figured prominently in arguments surrounding the debate.

Mental simulation of spatial events

Examples:
Searching a room in your "mind's eye",
mentally rotating an object, etc.

Review:
The analysis of this, and later, phenomena demands a review of the differences between the types of perceptual representation involved in the key frameworks. I take Barsalou's (in press) perceptual symbol systems, Johnson's image schemas (1987), and Kosslyn's mental imagery work (1976) as examples of the most detailed efforts at describing concept-like processing in perceptually-based systems.

In terms of the level of perceptual detail involved, Barsalou's perceptual symbols fall approximately halfway between Kosslyn's "picture-like" mental images and Johnson's highly abstract image schemas. Kosslyn's images are like percepts, in the sense that they have resolution, scale, etc. (recall, however, that Kosslyn does not claim that such a system can alone account for all conceptual activity). Barsalou's perceptual symbols are inextricably tied to a particular perceptual system, but are abstracted through a filtering process, and therefore not like conscious perceptual experiences. Instead of "seeing" stripes on a mental image of a tiger, Barsalou's framework instead describes the activation of a "stripedness" symbol in conjunction with a "tiger" symbol in the visual cortex. Johnson's image schemas are amodal, and even further abstracted. He focuses much of his argument on a few very general schemata, like "containment", which, he believes, underlie most human understanding. While he ultimately believes that all conceptual activity must be grounded in such sensorimotor representations, it's not clear how detailed the schemata can get, nor what additional mechanisms might be needed to provide a complete account of conceptual activity.

Discussion:
The following is a sample of some of the most compelling and influential empirical evidence relevant to mental imagery:

1. Kosslyn's "image tracing task" (1976) asked subjects to visualize a rabbit beside either an elephant or a fly. Subjects were then asked questions about detailed features on the rabbit's face. In the "fly" condition, in Kosslyn's interpretation, the rabbit "appeared" large in the mental image, and thus subjects were able to answer such questions quickly. In the "elephant" condition, subjects showed delays appropriate for having to "zoom in" on a picture frame in order to "see" the rabbit's features.

2. In mental rotation tasks with abstract figures (e.g., Shepard & Metzler, 1971), subjects show delays proportional to angle of rotation, as if they are rotating a "real" image spatially.

3. The V1 area of primary visual cortex has been shown to become active during mental imagery (Kosslyn et al., 1995).

In a domain in which compelling evidence is hard to come by, this is certainly strong evidence for some non-propositional, sensorimotor *component* to higher cognition. Proponents of propositional views argue that the experience of a mental image is a constructed by-product of some types of conceptual activity. However, they have trouble explaining the differential delays in the image tracing and mental rotation tasks: if the concept of "rabbit" has features associated propositionally, those features should be equally accessible, regardless of the context of the "scene".

Within the realm of perceptual explanations, the image tracing and mental rotation tasks appear to be evidence for a full imagery component (à la Kosslyn) rather than perceptual schemata (à la Barsalou), which abstract out issues of scale and resolution. Indeed, with respect to this evidence, the schematic approaches have no obvious advantage over the propositional ones. This leaves the perceptual approach with a dilemma. Abstraction is valuable in that it provides for "concept coherence", our experience of concepts as distinct, coherent entities. Also, it supports easier explanations of object identification (e.g., matching the concept of dog to the wide variety of perceptual conditions under which dogs are viewed) by abstracting away the details in which much of the perceptual variability lurks. However, by abstracting out issues like scale and resolution, the schematic theories need to reintroduce a mechanism by which perceptually detailed scenes are reconstructed to explain data like Kosslyn's. This leaves them with the same basic problem as the propositional theorists.

Furthermore, though this phenomenon of mental simulation does support a perceptual component of cognition, it does not necessarily imply that the entire conceptual system is sensorimotor-based. Kosslyn's dual-code (imagistic and propositional) system can also account for this evidence, though not without raising additional issues. The difficulty with such a system lies in giving an account of the apparently seamless interface between the two. Consider, for example, a task that combines perceptual and (classically) conceptual information, such as mentally searching your office space for all items worth more than $50. There is inevitably overlap in the kinds of entities each system can represent; it is a hard task indeed to explain how these redundant representations are kept consistent with each other, and no satisfactory account currently exists. I pursue this issue in greater depth, later in this section.

Using concepts that are heavily context-dependent and
Reasoning with ad hoc categories

Examples:

- context-dependent concepts: "normal", or "large" ("large molecule" vs. "small planet")
- ad hoc categories: items in the room that are likely to get stolen

Discussion:

These two phenomena raise many of the same issues: both stress the ability of the conceptual system to construct meaning dynamically. The term "ad hoc" refers to categories which are defined on the spot, but nevertheless are subject to fluid reasoning just like any more "permanent" category (e.g., "animal"). Barsalou, Johnson, and Clark & Grush all focus explicitly on concepts as representations that are dynamically constructed in working memory. This is in contrast to the traditional propositional approach, which characterizes the conceptual system as a (relatively) static structure in long term memory.

Context dependence has proven to be a fundamental characteristic of linguistic meaning. For all concepts, different aspects are highlighted under different circumstances. For example, the concept of "mammal" is often presented as an example of a well-defined "natural kind", and is thus presumably less subject to the effects of context. Nevertheless, using a term like "mammal" to describe my opponent in the context of a sporting event would likely to be taken to be derogatory. Would a fixed, propositional concept of "mammal" therefore contain every feature which might possibly be relevant under any circumstance? To take the argument further, consider the ad hoc category given as an example above. I can reason about the items in the room that are likely to get stolen as a unitary concept, but do they really explicitly share that feature in their long term representation? In other words, does every item in my possession carry a "likelihood- of-being-stolen" feature, as well as a "likelihood-of-being-destroyed-during-a-flood" feature, and so on? There would seem to be an infinite number of such features, if we don't restrict their level of detail.

So clearly, feature lists cannot tell us everything we need to know about concepts: concept meanings must have a dynamic component, presumably related to the reasoning processes in the context of which we generate understanding. Barsalou (1987) proposes such a dynamic view:

> Rather than being retrieved as static units from memory to represent categories, concepts originate in a highly flexible process that retrieves generic and episodic information from long-term memory to construct temporary concepts in working memory. Because this concept construction process is highly constrained by goals, context, and recent experience, the same concept is rarely if ever constructed for a category. (p. 101)

Barsalou's “simulations”, Clark & Grush's “emulations”, and Lakoff's “theory-based concepts” represent three different versions of this dynamic, constructivist view. Barsalou differentiates the long-term and short-term components by denoting the former as "simulators", the latter as "simulations". Simulators simply retain the core schemata in memory, but our experience and use of concepts comes only when various simulators are activated and combined in dynamic perceptual simulations. In Lakoff's version, the long term portion is the theory structure, and the active portion of a concept is a type of dynamic reasoning, based on those theories (e.g., a theory about what kinds of things thieves are likely to take, etc.). Clark & Grush's emulators, though not presented in the context of explaining high-level cognitive activity, do suggest that the dynamic emulation of basic sensorimotor patterns might be the core around which early symbols are constructed. Such emulations would necessarily incorporate context, as they originate under circumstances in which they must be able to continuously interact with the corresponding perceptual input.

While it is clear that the context sensitivity of concepts is commonly a component of embodied, perceptual theories, we need to consider whether they are really any better at explaining it. Certainly, neither perceptual nor propositional systems *preclude* the possibility of taking context into account. The argument made by the embodied community is that the perceptual system can't *help* but take context into account. This may be true during the initial creation of a perceptual symbol (in fact, explaining how the "core" schemata are *extracted* from the context is problematic), but once a symbol system is in place, and symbols are being reactivated to construct a conceptual scene, it's not clear how perceptual symbols have any more claim to context-dependence than abstract, conceptual symbols. Barsalou's account of a perceptual symbol for a concept like "larger" is as atomic as any: it simply binds two other perceptual symbols as arguments, as if it were a propositional system. Other perceptual theorists (e.g., Johnson) hint at the increased flexibility that a perceptual representation provides, but offer no detailed explanation. Context sensitivity may be impossible to describe in the brittle, definitional style of propositional system, but this doesn't necessarily hold for fuzzy or probabilistic forms of abstract concept systems. Basically, there is currently no account of graded concepts in *either* system that clearly outperforms the other. I would suggest the claims regarding the handling of dynamic concepts arise more from a bias which reflects the focus of the research in each approach, rather than the actual capabilities of the representations.

Understanding, reasoning with, and extending words without clear definitions

Examples:

We make sense of the sentences "Greg opened the book" and "Greg opened the store", but not "Greg opened the tree". How do we define "open"? What is an acceptable object to "open"? Furthermore, note that "opened the tree" *could* make sense in a very limited context, like "tree surgery".

If I understand the phrase "open the coconut", I can also extend my understanding of the term "open" to understand the phrase "open the walnut", even if I have never seen a walnut opened.

Discussion:

The first example emphasizes two facts: (1) concepts are not always easy to explicitly define, and (2) meaningfulness is a continuum, a phenomenon which recalls the previously discussed issues of graded category membership and context-dependence. With regard to the first point, Barsalou (1987) cites evidence that many very common terms are surprisingly difficult for people to define. Again, these kinds of examples are a real problem for propositional systems which suggest that concepts are defined by fixed sets of necessary and sufficient conditions, or invariant feature lists. Such systems seem too rigid to support the obviously fuzzy, unstable nature of concepts.

The perceptual-schematic explanation for this example contends that any object which can be successfully matched against a "container" or "closed form" schema can be understood as the object of "open". The "definition" of open itself is a dynamic schema which actually represents the opening of a closed form schema. The term never needs precise definition in any other form. Presumably, "tree" is a poor object for "open" because it provides a poor perceptual match (unless the "knife-wielding tree surgeon" completes the picture, in which case it is marginally acceptable). While this description is intuitively appealing, it relies on a perceptual matching process which has not been adequately defined in *any* theory to date. It is not self-evident that extracting and/or comparing some core perceptual "dogness" from the incredible range of perceptual traces which invoke the idea "dog" is any easier than the extracting the corresponding abstract symbol. Nevertheless, perceptual pattern matching is clearly something that the human brain is good at, so proponents of the perceptual view at least have some indirect evidence that they may yet find the mechanism(s) they need.

Now let us reconsider this phenomenon of fuzzy definitions from the perspective of a non-perceptual symbol system. As established in the previous section, a list of discrete features complete enough to represent every possible characteristic (e.g. "book: can_be_construed_as_container") is implausible. There must be a dynamic, inferential step involved in the construction of meaning that establishes these kinds of characteristics; it is the nature of this step that is therefore central to our debate.

One possible propositional explanation for such "fuzzy" language understanding would be exemplar based: you don't understand "opened the tree" because you've never heard "open" used in that context before. Nevertheless, it's clear that we *can* easily extend the meanings of words in novel, but systematically related ways. The second example above is an instance of this phenomenon: "open the walnut" makes sense even though I have no obvious definition to apply to this novel circumstance.

Mark Johnson's (1987) answer to this second example is that systematic meaning comes from metaphor. Johnson believes that systematic, conventional metaphorical systems underlie nearly every domain of human understanding. A classic example he explores in detail is that of theories (an abstract concept) understood in terms of buildings (a concrete one). Theories have support, foundations, form, strength, framework, etc. Thus, meaning is grounded in sensorimotor experience, and nearly all such grounding occurs in the context of metaphor. The strongest evidence for this view, Johnson claims, is linguistic. First, the metaphors show up in the language we consistently use to describe concepts. Second, he argues that the different meanings of polysemous words can appear to be very different on the surface, but can almost always be related by some abstract, metaphorically-extended image schema. Metaphor thus provides a high-level explanation for why some meanings are better than others (better and worse matchings between the domains), how we can come to accept new uses for a word, and how we can ground the meaning of more abstract concepts in more directly physical ones.

To relate this explanation back to our debate, we need to consider whether this metaphorical system must be implemented in perceptual schemata, as opposed to a propositional format. Ostensibly, the answer is "no": the details of metaphorical matching are not provided in such great depth as to preclude a probabilistic, propositional implementation. Nevertheless, Johnson argues the answer is "yes", and cites the symbol grounding problem as the reason. Ultimately, he argues, the meaning of abstract symbols must be grounded in something other than other abstract symbols; they must be grounded in experience. Since the sensorimotor system is most closely connected to experience, an appeal to parsimony draws Johnson to the conclusion that it must be *perceptual* schemata which underlie concepts. I consider this argument in greater detail with respect to the next phenomenon: abstract concepts.

Reasoning with abstract concepts

Example:
Comparing the concepts of "justice" and "fairness"

Discussion:

The concept of "justice" has no obvious direct physical instantiation which might serve as the core of a perceptual symbol. Barsalou, who gives the most detailed account of perceptual schema formation, describes it as a process of abstracting/filtering to extract a subset of the non-conscious neural activity associated with the related percept. The problem, in this case, is that there is no obvious percept to abstract from. Clearly, there are associated ideas like balance, which plausibly have perceptual instantiations. However, Barsalou's framework demands that *all* aspects of the concept are represented in this format, which means being able to distinguish the subtle connotations differentiating, for instance, "justice" and "fairness". Barsalou (in press) attempts to explain abstract concept handling through a few examples, but the examples are carefully chosen, and he admits that the explanations are weak in their current form.

A similar argument can be leveled against Johnson's (even more abstract) image schemas, though he concedes amodal content and is less specific about whether additional mechanisms are needed to cover the whole range of conceptual activity. It is intuitively possible that the metaphor-based system is able to create enough associations to properly capture the subtleties Barsalou's system seems to lack, but there is no strong evidence to support this claim.

The argument from the propositional viewpoint is that the abstract/concrete dimension is almost irrelevant. Neither is harder to generate or explain than the other. This claim, however, fails to acknowledge the issues surrounding the symbol grounding problem. To say that it is easy to bring a new abstract symbol into existence and attach the label "justice" to it is one thing; to provide an account for how the meaning is truly established is quite another. If associations to other abstract symbols are not enough to account for meaning, and thus an experiential, perceptual component of the conceptual system is actually necessary, then at best, the propositional account is no better off than Johnson's metaphor-based account. In fact, it may be worse, as it may need a more sophisticated explanation of the perceptual-to-conceptual transduction process than the image schema framework, as image schemas are already inherently perceptual. Since some filtering or transduction is needed in both cases, however, it's not clear how strong an argument this really is.

In short, the propositional side of the debate typically claims victory over this issue of abstract concepts, but on closer inspection it appears to be a difficult phenomenon for *anyone* to explain.

Reasoning about impossible events

Example:
Reasoning that dragons weigh less on the moon than on earth.

Discussion:

This kind of example is precisely what the propositional systems used in traditional AI are good at. Although dragons are mythical, they are mythical *animals*, and thus inherit the default features of animals. In turn, animals inherit the feature of "mass" from physical objects, and the fact that objects with mass weigh less on the moon can be conveniently captured in a rule. Of course, in a disembodied system, the "reality" of the situation described is, in some sense, irrelevant. The world consists of the contents of the propositions, and whether the propositions are labeled "dragon", "huju", or "rock", their ontological status can be reduced to simply another recorded feature.

From the embodied view, ontological status is considerably more salient. Nevertheless, a perceptual understanding of the dragon sentence is not problematic, as one's experience with pictures of dragons could easily suffice to create a perceptual symbol like that of any other animal. Barsalou stresses the *productive* nature of the perceptual system: symbols may be recombined in an infinite number of ways. Inheritance is captured in the process of selecting what kinds of simulators are suitable to be included in what kinds of simulations: a "dragon" can participate in a simulation suitable for "animal", or indeed, "object with mass". There is no detailed account of the process of storing or determining what constitutes an acceptable participant in a particular simulation, although it presumably is based in perceptual experience.

As a footnote, it is interesting to ponder a potentially deeper issue, concerning *completely* novel physical entities, like "a new color, unlike any existing color", or a situation in 5 spatial dimensions. In theory, a completely abstract propositional system with no grounding in experience should have no more difficulty describing a 5-dimensional event than a 2-dimensional one. The former, though, seems distinctly more difficult for human beings to conceive of.

Integrating conceptual activity and physical activity

Example:
Finding the most "expensive", or "dangerous", item in a room

Discussion:

There is considerable evidence that physical activity interacts with many aspects of conceptual activity. Van den Bergh, Vrana, and Eelen (1990) found that when subjects were forced to indicate a preference between sets of visually presented letter pairs, subjects with typing experience showed a clear preference for pairs that were easier to type, whereas non-typists showed no clear preference. Glenberg (1997) cites evidence that memory for commands that were acted out is better than memory for the verbal description of the commands. Finally, Klatzky, Pellegrino, McCloskey, and Doherty (1989) demonstrated that subjects were able to make semantic sensibility judgements of phrases describing meaningful and nonsensical physical acts (e.g. "aim a dart" or "close a nail") faster when the subject was performing appropriate physical hand motions (e.g. pinching or clenching).

The embodied, situated view holds that interaction between the physical and conceptual is the *default* situation in human cognition. It is argued that perceptual symbol representations can be matched against perceptual/motor experience more easily than performing a percept-to-concept transduction. However, the details of such a matching system (as argued above) are so seriously lacking, that it's not clear that matching perceptual schemata is necessarily easier than extracting and matching propositional representations. "Expensive item", for instance, is an abstract concept; it's not obvious what an appropriate perceptual schema is, nor how it could be perceptually matched against a visual scene. Once again, it seems the data that nobody is able to account for greatly outweighs the advantages any one representation claims over another.

Before moving on, however, I would like to consider two additional pieces of experimental work which relate sensorimotor experience to conceptual activity, and provide some context for their possible integration. First, Rodney Brooks' work on embodied, situated mobots (described in section 4.2.7) has shown that, for some tasks at least, action in the world can be performed more efficiently (in less time, and with less computational power) if internal modeling and planning are minimized. As an example, Brooks (1991) describes a DARPA project aimed at building an autonomous vehicle that could drive down a road given visual input from the driver's viewpoint. The first attempts followed the classic sense-model-plan-act methodology, attempting to build a 3-dimensional internal model of the scene which a planning system would then manipulate to produce the appropriate output. This attempt failed because extracting the model proved too difficult, and the computation took too long. Two other teams, working independently, subsequently came up with successful solutions that worked directly from the 2-D picture input to continuously update the servo output, without additional modeling or planning. The system was thus able to run in real time, and the difficult problem of 3-D modeling was averted. One could argue that evolution would likely take the same shortcuts wherever possible, avoiding transduction to non-sensorimotor representations.

The second relevant piece of experimental work comes from Clark and Grush (in press), who argue that emulation of sensorimotor activity is already necessary to explain some motor phenomena, and offer it as a possible explanation of the evolutionary origin of conceptual phenomena as well. They present the case of smooth grasping motion, the study of which has shown that proprioceptive feedback is not fast enough to account for observed performance. To explain how this "impossible" level of performance is achieved, they describe an emulation system which dynamically models the system state faster than the real world input is received, thus "predicting" the future state of the system and responding appropriately. They then proceed to construct an elegant description of how such a system might become the evolutionary foundation of mental representations. They describe how sensorimotor emulation concurrent with activity would be an easy phenomenon to explain neurally, and how such emulation may have gradually become desynchronized from the external stimuli in order to provide these "predictions" which result in greater performance. The next step beyond desynchronization is complete detachment from sensorimotor stimuli, in which the organism can simulate the activity without performing it. This, too, is evolutionarily sound, in that it provides "practice" for making smooth, delicate motions. Such offline sensorimotor emulations could then form the foundation of perceptual symbols which could become (at least part of) a conceptual system.

What we get from these two examples is a sort of "proof of existence" that pure sensorimotor activity could possibly account for low level conceptual (or at least concept-like) behaviors. Unlike most of the other arguments we have considered, these descriptions are fairly complete, and testable. Of course, neither attempts to address the high-level concepts that have been the primary focus of this paper. Nevertheless, they at least suggest an approach that may be able to establish quite decisively whether the theory of embodied, perceptual cognition is constructed on a firm foundation.

Rule-based categorization

Example:
Use of a rule to categorize 323,456,391 as an odd number, despite never having seen it.

Discussion:
Rule-based categorization could be described as a case of quintessential propositional processing. Even though not all propositional theories rely exclusively on rule-based processing, all such systems (by definition) do *permit* the existence of explicit rules.

Advocates of the perceptual view argue that there's no need for "rule-like" behavior to be driven by explicit rules; countless examples of neural networks have shown this to be true (Rumelhart & McClelland, 1986). A perceptual simulator could easily capture the feature "last digit is 1,3,5,7 or 9" (all of the components of which have perceptual interpretations).

The only problem for Barsalou's framework, in particular, is explaining the fact that the same rule must be redundantly stored in both the visual and auditory system. Barsalou's insistence on placing all content within particular modalities demands that the concept of "oddness" be separately represented for each. This seems to be a very serious, and unnecessary, breach of the rule of parsimony. Johnson's acceptance of the multimodal and amodal (though percept-like) nature of some conceptual content avoids this implausible consequence.

Subtle perceptual categorization

Example:
Distinguishing male and female faces, without necessarily knowing how we do it.

Discussion:
This issue is partially a question of what's accessible to consciousness. It is certainly *possible* that there is a very complex set of criteria that a propositional system could apply probabilistically (and non-consciously) to differentiate male and female faces. However, the "hard work" of detecting and extracting such subtle features is clearly being done in the perceptual system (or in the percept-to-concept transducer). Thus, even a propositional explanation must admit that the purely abstract part of the system is doing very little of the work in this example.

As for the schemata-based systems, they save having to describe the additional transduction step, but are really no closer to describing the mechanism through which such subtle categorization is performed. They lack a principled explanation of how some of the abstracted features (like "stripedness") are accessible while others (like "maleness") are not, especially since Barsalou ascribes the task of feature abstraction to the (largely conscious) attentional system.

Full image-based systems like Kosslyn's might have a better explanation through prototyping/exemplar theories in which *examples* of male and female faces ultimately underlie the ability to differentiate. Such a process could be inaccessible to consciousness, and free to derive unlimited complex feature combinations. Ultimately, however, the nature of a general, perceptual similarity comparator has proven so difficult to characterize that it's unclear whether the exemplar approach is actually any closer to an answer. Once again, we have a case where the data is apparently difficult for *all* sides to explain.

One additional piece of evidence relevant to the issue of categorization comes from Glaser's (1992) review of some reaction time experiments which manipulated stimulus modalities. In particular, subjects in one experiment were able to name the superordinate category of a stimulus faster if the stimulus was presented as a picture rather than as a word. While this experimental design was plagued, in my opinion, by a number of possible confounds (for instance, the verbal response might have been delayed by interference with the stimulus in the word case), the results suggest that image processing is not considerably slower than word processing, which is what might be expected in a propositional system.

Development

Examples:

Phylogenetic development; how we developed smoothly from, presumably, purely perceptual ("stimulus-response") creatures

Ontogenetic development; how (and whether) we develop from concept-free or concept-limited infants to highly conceptual adults

Discussion

In section 4, I provided a very brief outline of the theories of concept development proposed by Piaget and Mandler. Piaget assumed that the perceptual system precedes a separate conceptual system in infant development. Mandler, however, cites considerable evidence of symbolic behaviors in infants less than 10 months of age (e.g., Baillargeon, De Vos, & Graber, 1989, shows recall by infants as early as 8 months of age), and concludes that the conceptual activity co-evolves along with perceptual activity. This might be interpreted as indirect support of close ties between perceptual and conceptual activity, though Mandler herself nevertheless assumes (without directly arguing the point) that conceptual activity is the responsibility of a distinct system. She admits having no evidence suitable for explaining the interface between perceptual and conceptual activity.

As described above, Clark & Grush (in press) provide a particularly elegant and plausible description of the phylogenetic origins of mental representations, which might just as easily apply to ontogeny as well. While their description places the *first* true mental representations in the sensorimotor system, this of course does not preclude the possibility of the subsequent creation of a distinct conceptual system. Any such account, however, would need to explain what evolutionary drive caused the creation of this separate system. As yet, no one has offered such an explanation.

With respect to development, then, I suggest that early symbolic thought seems likely to be based in, or at least strongly influenced by, sensorimotor activity. This provides some (albeit weak) evidence in support of experiential grounding, and makes the job of describing a distinct conceptual system somewhat harder.

6. Conclusions

The analysis above has, out of necessity, been limited in scope. Nevertheless, consideration of the major arguments and evidence from each of these domains does lead to several conclusions. Perhaps the most obvious conclusion is that in *every* aspect of this debate, the number of unanswered questions so greatly outweighs the available evidence, that it is perhaps premature to expect a more thorough analysis to produce more definite answers.

Despite these difficulties, I believe several reliable conclusions have come to light in the course of this analysis:

- Many of the phenomena, despite being presented as arguments for one side or the other in the literature, are in fact difficult for *any* current framework to explain. These arguments often appeal to an intuitive sense of the "naturalness" of a representation, but their advantages are not further elaborated, and they do not constitute decisive evidence. For instance, rule-based categorization is presented is support of propositional representation, when it presents no deeper challenge for perceptual frameworks. On the other hand, context sensitivity of concepts is presented as supporting evidence for perceptually-oriented systems, when some propositional approaches seem no farther from a solution.

- There is considerable evidence that perceptual representations play *some* role in higher cognition. At the very least, it seems misguided to expect to describe the human conceptual system without taking the nature of the sensorimotor system into account. Even if the conceptual and perceptual systems *are* separable, one would expect the nature of the perceptual system to profoundly influence the structure of the conceptual system.

- The evidence from the image tracing task and mental rotation tasks seem to support the conclusion that considerable perceptual detail (e.g. realistic spatial orientation and scale) is potentially available in mental imagery, a fact which may be troublesome even for the current generation of perceptual schematic theories. Clearly, different levels of abstraction come with different tradeoffs, and it appears that no single level is capable of providing all of the answers.

- The argument for entirely modal systems seems to introduce unnecessary restrictions and challenges that a combination modal/amodal system need not face. Explaining multimodal phenomena without unnecessary duplication of representations, and without placing *any* conceptual content outside of modality-specific brain regions, seems implausible and unnecessary. There is little reason to believe that reduced versions of direct perceptual input (a la Barsalou) can truly capture all the nuances of meaning, and represent abstract concepts that have no obvious perceptual correlate. By allowing amodal conceptual content, percept-like representations (e.g. image schemas) can bridge modalities and engage in other forms of associations (i.e. analogistic) without giving up the considerable advantages of keeping representations "close to the sensorimotor surface".

- Propositional systems which describe concepts as static structures defined by necessary and sufficient conditions, or by invariant feature lists, are doomed to fail. Overwhelming evidence suggests that concepts are dynamic, fuzzy, and highly context-dependent.

- In light of the above, an atomistic view of concepts as clearly individuated symbols may be in jeopardy as well, although the evidence is not overwhelming. It should be noted, however, that perceptual accounts like Barsalou's are just as subject to this criticism as many propositional systems.

- The symbol grounding problem constitutes a considerable obstacle for any propositional systems which reject a close relationship to perceptual representations. There are simply no convincing arguments for how to endow a conceptual system with real meaning without invoking a tight coupling between concepts and the experience of interacting with the world.

In light of these conclusions, how do we decide whether it is useful and meaningful for cognitive science to distinguish between concepts and percepts? I would argue that both modal and amodal representations directly play the role of conceptual content: there are some things which seem too abstract to describe compactly in modal form, and some things which are so easily described in modal form that it would be a breach of parsimony to translate the content to another representation.

It may, however, be more appropriate to de-emphasize the impact of such a statement. "Amodalness" comes in degrees; there is a considerable range of possible "distances" from the sensory apparatus. To expect to be able to draw a hard line between "modal" and "amodal", let alone "perceptual" and "conceptual", might be an oversimplification of matters. To posit a "dual-representation" system suggests two distinct representational formats with a definite transduction mechanism between them. This view demands an explanation of the nature of the transduction, as well as the developmental origins of the distinct conceptual system. Perhaps this is not the appropriate view. The evidence suggests a wide range of representations, at a wide range of abstraction. It seems plausible that this range might represent a continuum, in which representations retain as much of their perceptual character as possible for the task at hand. Multiple representations of the same entity may be active simultaneously, tightly coupled but filtered to permit different kinds of manipulation.

I have concluded that it is highly unlikely that there is *not* a perceptual component to higher cognition. The evidence is not as clear that it is *necessary* to have a propositional component that can be interpreted as the rule-based manipulation of unambiguous, discrete symbols. I have suggested that even if such a system exists, it would still need to be grounded in the meaning available in the perceptual/motor system. Together, this evidence does seem to suggest a direction: perhaps an increased appreciation for embodiment and the role of the perceptual system *is* appropriate, if only for pragmatic reasons. Such an approach could be conservative, in the sense of beginning with reliable, grounded perceptual properties and positing more sophisticated machinery only when it is needed.

Let me emphasize that I do *not* arrive at this conclusion because I have been so impressed by the explanatory successes of the existing, perceptually-oriented frameworks. Instead, it is the appeal of the "firm ground", the potential to actually approach a complete description of the low-level foundations of symbolic thought through work like that of Brooks, Beers, and Clark & Grush. To describe the complete human conceptual system from the ground up is, to put it mildly, a Herculean task. Nevertheless, the degree to which high-level concept work is based in intuition and speculation is disturbing. The advantage of the embodied, perceptual approach may be that it provides a concrete target, a low-level context in which to situate and target the study of the high-level human conceptual system.

References

Baillargeon, R., DeVos, J., & Graber, M. (1989). Location memory in 8-month-old infants in a non-search AB task: Further evidence. *Cognitive Development*, 4, 345-367.

Barsalou, L.W. (in press). Perceptual symbol systems. *Behavioral & Brain Sciences*.

Barsalou, L.W. (1987). The instability of graded structure: Implications for the nature of concepts. In Neisser, U. (Ed.) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press.

Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173-215.

Brooks, R. (1991). Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Aritificial Intelligence*. Morgan Kauffman.

Clark, A. & Grush, R. (in press). Towards a cognitive robotics. *Adaptive Behavior*.

Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.

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

Fodor, J.A. & Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Glaser, W.R. (1992). Picture naming. *Cognition*, 42, 61-105.

Glenberg, A.M. (1997). What memory is for. *Behavioral & Brain Sciences*, Mar. 20(1), 1-55.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335-346.

Harnad, S. (1987). *Categorical Perception: The Groundwork of Cognition*. Cambridge: Cambridge University Press.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Imagination, Reason and Meaning*. Chicago: University of Chicago Press.

Kirsh, D., & Maglio, P. (1992). Some epistemic benefits of action: Tetris, a case study. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Klatzky, R.L., Pellegrino, J.W., McCloskey, B.P., & Doherty, S. (1989). Can you squeeze a tomato? The role of motor representations in semantic sensibility judgements. *Journal of Memory and Language*, 28, 56-77.

Kosslyn, S.M., Thompson, W.L., Kim, I.J., & Alpert, N.M. (1995). Topographical representations of mental images in primary visual cortex. *Nature*, 378, 496-498.

Kosslyn, S.M. (1976). Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval time. *Memory & Cognition*, 4, 291-297.

Lakoff, G. (1987). Cognitive models and prototype theory. In Neisser, U. (Ed.) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press.

Lenat, D.B. & Guha, R.V. (1990). *Building Large Knowledge-Based Systems*. Reading, MA: Addison-Welsley.

Mandler, J.M. (1997). Development of categorization: Perceptual and conceptual categories. In Bremner, G., Slater, A., & Butterworth, G. (Eds.) *Infant development: Recent advances*. Hillsdale, NJ: Erlbaum.

Mandler, J.M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review*, 99(4), 587-604.

Neisser, U. (1987). From direct perception to conceptual structure. In *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, Neisser, U. (Ed.). Cambridge: Cambridge University Press.

Newell, A. and Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3), 113-126.

Piaget, J. (1952). *The Origins of Intelligence in Children*. New York: International University Press.

Port, R. & van Gelder, T. (1995*). Mind as Motion: Dynamical Perspectives on Behavior and Cognition*. Cambridge, MA: MIT Press.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B.B. Lloyd (Eds.) *Cognition and Categorization*. Hillsdale, NJ: Erlbaum.

Rumelhart, D.E., & McClelland, J.L., Eds. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 1. Foundations*. Cambridge, MA: MIT Press.

Shepard, R. & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701-703.

Thompson, E. (1995). *Colour vision : a study in cognitive science and the philosophy of perception.* London: Routledge.

Van den Bergh, O., Vrana, S., & Eelan, P. (1990). Letters from the heart: Affective categorization of letter combinations in typists and nontypists. *Journal of Experimental Cognition: Learning, Memory, & Cognition*, 16, 1153-1161.

Vera, A. & Simon, H. (1993). Situated action: A symbolic interpretation. *Cognitive Science*, 17, 4-48.