THE ROYAL SOCIETY

Published online 22 May 2003

Abstraction in perceptual symbol systems

Lawrence W. Barsalou

Department of Psychology, Emory University, Atlanta, GA 30322, USA
(barsalou@emory.edu, http://userwww.service.emory.edu/~barsalou/)

After reviewing six senses of abstraction, this article focuses on abstractions that take the form of summary representations. Three central properties of these abstractions are established: (i) type–token interpretation; (ii) structured representation; and (iii) dynamic realization. Traditional theories of representation handle interpretation and structure well but are not sufficiently dynamical. Conversely, connectionist theories are exquisitely dynamic but have problems with structure. Perceptual symbol systems offer an approach that implements all three properties naturally. Within this framework, a loose collection of property and relation simulators develops to represent abstractions. Type–token interpretation results from binding a property simulator to a region of a perceived or simulated category member. Structured representation results from binding a configuration of property and relation simulators to multiple regions in an integrated manner. Dynamic realization results from applying different subsets of property and relation simulators to category members on different occasions. From this standpoint, there are no permanent or complete abstractions of a category in memory. Instead, abstraction is the skill to construct temporary online interpretations of a category's members. Although an infinite number of abstractions are possible, attractors develop for habitual approaches to interpretation. This approach provides new ways of thinking about abstraction phenomena in categorization, inference, background knowledge and learning.

Keywords: abstraction; concept; interpretation; dynamic; simulation; embodiment

1. INTRODUCTION

Abstraction is a central construct in cognitive science. Rather than just having one sense, abstraction has at least six, as follows.

(i) *Categorical knowledge*. Abstraction can simply mean that knowledge of a category has been abstracted from experience, such as abstracting the category of *CHAIRS* from the settings in which they occur. (Italics will be used to indicate concepts, and quotes will be used to indicate linguistic forms (words, sentences). Thus, *CHAIR* indicates a concept, and 'chair' indicates the corresponding word. Within concepts, uppercase words will represent categories, whereas lowercase words will represent properties of categories (e.g. *CHAIR* versus *seat*) and relations between properties (e.g. above for the relation of the *CHAIR*'s back to its *seat*).) Nearly all accounts of knowledge are comfortable with this sense, including rule-based, prototype, exemplar, connectionist and embodied theories.

(ii) *The behavioural ability to generalize across category members*. Another uncontroversial sense of abstraction is that people can summarize the properties of a category's members *behaviourally*. All theories agree that people state generics, such as 'cats have fur,' and quantifications, such as 'some mammals swim'. Behaviourally, people produce abstractions.

(iii) *Summary representation*. The cognitive bases of the behavioural abstractions in sense (ii) are controversial. In some theories, behavioural abstractions reflect underlying summary representations of category members in long-term memory. According to these views, when people generalize behaviourally, they describe an underlying summary representation, such as a declarative rule, a statistical prototype or a connectionist attractor. Importantly, the summary representations in sense (iii) are not necessary to produce the behavioural abstractions in sense (ii). For example, exemplar models store only exemplars in memory—not summary representations—and produce behavioural abstractions by scanning and summarizing exemplars online (e.g. Hintzman 1986).

(iv) *Schematic representation*. A second controversial sense is that schematic representations represent categories in memory, where 'schematic' refers to summary representations being sparser than exemplars. Thus a schematic representation might abstract the critical properties of a category's exemplars and discard the irrelevant properties (e.g. the geons of Biederman 1987). Also, properties in a summary representation may be distorted to idealize or caricature a category, helping to distinguish the category from others (e.g. Posner & Keele 1968; Rhodes *et al.* 1987; also see Barsalou 1985; Palmeri & Nosofsky 2001).

One contribution of 16 to a Theme Issue 'The abstraction paths: from experience to concept'.

DOI 10.1098/rstb.2003.1319

(v) *Flexible representation*. A third controversial sense of abstraction is that summary representations can be applied flexibly to many different tasks, including categorization, inference, language comprehension, reasoning, etc. From this perspective, increasing abstractness allows a representation to become increasingly flexible (e.g. Winograd 1975).

(vi) *Abstract concepts*. Finally, abstraction can refer to the abstractness of concepts, ranging from concrete (e.g. *HAT*) to abstract (e.g. *COURAGE*). When concepts become detached from physical entities and more associated with mental events, they become increasingly abstract (e.g. Paivio 1986; Barsalou 1999; Wiemer-Hastings *et al.* 2001).

Although all six senses of abstraction are important, the focus here will be on one of its more controversial ones: abstraction as summary representation (sense (iii) above). In what follows, 'abstraction' will refer solely to this sense. The goal of this article will be to develop an account of summary representations within the framework of perceptual symbol systems (Barsalou 1999). For an extended version of this article, see Barsalou (2004).

2. THREE PROPERTIES OF SUMMARY REPRESENTATIONS

Three properties are central to abstractions that take the form of summary representations: type–token interpretation, structured representation and dynamic realization.

(a) *Property 1: type–token interpretation*

Pylyshyn (1973) proposed that cognition is inherently an interpretive process. In the debate on mental imagery, he argued that cognitive representations are not like the holistic bit-mapped recordings in cameras, video recorders and audio recorders. Many other perception researchers would agree (e.g. Hochberg 1998). Instead, Pylyshyn argued, cognitive representations are interpretations of experience. To construct an interpretation, concepts in memory type the components of sensory-motor experience to produce type–token propositions. On walking into an office, for example, the concepts for *COMPUTER*, *TABLE* and *LAMP* become bound to particular objects, thereby creating type–token propositions of the sort *COMPUTER*(object-89), *TABLE*(object-23), etc. Such propositions essentially make claims about the world that can be true or false, such as the belief that object-89 is a *COMPUTER* (e.g. Church 1956).

A component of experience can be interpreted in infinite ways. For example, object-89 could be interpreted alternatively as *ARTEFACT*(object-89), *OFFICE EQUIPMENT* (object-89), *ELECTRONIC DEVICE*(object-89), *DEVICE THAT REVOLUTIONIZED THE MODERN WORKPLACE*(object-89) and so forth. An infinite number of true interpretations of an individual exist and also an infinite number of false interpretations, with each interpretation providing a different perspective on the object.

Once a type–token proposition exists to interpret an entity or event, the proposition provides extensive inferential knowledge. Once something is interpreted as a *COMPUTER*, inferences follow, such as that it requires electricity, can be used for e-mail, is easily breakable and so forth. If the object were interpreted instead as *SOMETHING THAT THIEVES STEAL*, different inferences would follow (e.g. the computer should be locked to its table). In all cases, these inferences constitute further propositions that become linked to the type–token mappings that triggered them.

From this standpoint, propositions underlie representations of the world, not bit-mapped recordings (see Haugeland 1991; Dretske 1995; Barsalou 1999). A representation of a computer is not a holistic recording, but a set of propositions that interpret it. Most importantly for the purpose of this paper, Pylyshyn assumes that abstractions underlie the interpretive process. The types in type–token propositions are abstractions for properties, objects, events, relations and so forth. After a concept has been abstracted from experience, its summary representation supports the later interpretation of experience. Therefore, abstractions (as summary representations) underlie interpretation.

(b) *Property 2: structured representation*

Concepts do not typically interpret experience individually but are organized into structured representations that establish relations between individual type–token propositions. Rather than *COMPUTER*(object-89) and *TABLE*(object-23) being independent, a spatial concept, such as *on*, might organize them into a structured proposition, such as

on(*upper-region* = *COMPUTER*(object-89),
lower-region = *TABLE*(object-23)).

Considerable empirical evidence indicates that structured representations pervade human knowledge. Some of the strongest evidence comes from work on concepts and categorization, where researchers have found robust evidence for relational structure in experiments designed to detect it (e.g. Goldstone & Medin 1994; Markman & Gentner 1997; also Barsalou 1992; Barsalou & Hale 1993). Categorizing exemplars, judging their similarity and drawing categorical inferences all rely heavily on structured relations—not only on independent properties. Additionally, the process of combining individual concepts into structured representations underlies the process of conceptual combination (e.g. Rips 1995; Hampton 1997; Wisniewski 1997).

Further evidence comes from analogy, where structured representations are clearly implicated in people's ability to extend relational systems from one domain to another (e.g. Gentner & Markman 1997; Holyoak & Thagard 1997). Similar evidence comes from language comprehension, where complex propositional structures underlie the meanings of texts (e.g. Kintsch & van Dijk 1978; Graesser *et al.* 1994). Finally, various theorists have argued that structured representations are a hallmark of human cognition, which any theory must explain (e.g. Fodor & Pylyshyn 1988).

For these reasons, a second fundamental property of abstractions is their participation in complex interpretive systems. Abstractions do not simply interpret isolated components of experience, but assemble into structured representations that interpret complex structure in the world.

(c) *Property 3: dynamic realization*

The abstractions that represent a category are notoriously difficult to specify. In attempting to specify the abstractions that represent particular categories, it is usually impossible to specify them fully (e.g. Barsalou 1993). Artificial intelligence researchers often experience similar difficulty in articulating abstractions when programming knowledge into intelligent systems. In general, three problems arise when attempting to specify the abstractions that underlie a category.

(i) *Identifiability*. What information about a category should an abstraction include? For example, what abstractions should be used to represent restaurant visits (Schank & Abelson 1977)? Of everything that could possibly occur in these visits, what should a summary representation of them include? Only the properties that are relatively invariant across restaurant visits? What about properties that are true occasionally? How should differences between individuals and cultures be handled? How does a theorist determine when an abstraction is complete? It is a daunting task to specify the content of an abstraction completely.

(ii) *Justification*. How does a theorist justify the inclusion of particular information in an abstraction? In artificial intelligence, knowledge engineers often select abstractions intuitively that best serve a specific application that they are trying to develop. Problematically, however, no clear principles exist for justifying the inclusion of particular information in abstractions.

(iii) *Rigidity*. Typically, exceptions arise for an abstraction (e.g. Wittgenstein 1953). A frequent criticism of the restaurant script of Schank & Abelson (1977) was that it did not cover unexpected deviations and unusual restaurant visits. In response, Schank and Abelson suggested that different *tracks* through a script handle special cases. Problematically, however, infinitely many tracks are required to handle all possibilities. Furthermore, how do people process novel cases, which they often handle relatively effortlessly (e.g. visiting a new type of restaurant)? To date, we have no satisfactory account of how abstractions can handle such variability.

The identifiability, motivation and rigidity problems could be viewed as indicating that we simply need a better methodology for discovering abstractions. Alternatively, however, there may be no correct abstractions to discover. Instead of a single abstraction representing a category, an infinite number of abstractions may be constructed online to represent a category temporarily (Barsalou 1987, 1989, 1993). If this latter conclusion is correct, then studying the *skill* to construct temporary abstractions dynamically may be more informative scientifically than attempting to discover one particular abstraction that represents a category. For this reason, I assume that a third important property of abstractions is their dynamic realization. In the literatures that address abstraction, this is not a standard assumption.

(d) *Existing theories of summary representations*

Classic theories of representation typically do a good job of implementing the first two properties just discussed for abstractions: type–token interpretation and structured representations. Good examples of these theories and how they implement these two properties can be found in Winograd (1972), Newell & Simon (1972), Schank & Colby (1973), Bobrow & Collins (1975), Collins & Loftus (1975), Anderson (1976), Schank & Abelson (1977) and Charniak & McDermott (1985) (also see the articles by Zucker (2003) and Holte & Choueiry (2003)). Problematically, however, these theories are poor at handling the third property of abstractions, namely, dynamic realization. As just reviewed, these theories typically have difficulty in identifying and justifying the content of abstractions, and they produce abstractions that are overly rigid.

By contrast, connectionist approaches to abstraction succeed beautifully at handling the flexibility of abstractions and also at implementing simple type–token interpretation. Where connectionist theories have difficulty is in implementing structured representations. Although various proposals have been suggested, none has convinced many researchers that it provides a psychologically plausible account. Barsalou (2004) presents a more detailed discussion of these various theories, along with their strengths and weaknesses.

This article's theme is that a satisfactory and powerful account of abstraction can be developed within the framework of perceptual symbol systems (Barsalou 1999). This theory implements structured interpretation naturally and elegantly, while simultaneously implementing dynamic realization.

The next section lays the groundwork for this approach to abstraction. The final two sections then illustrate how the various properties of abstraction arise within this framework. The goal here is to outline the basic architecture of this approach, and to show how it implements a satisfactory account of abstraction. Although not provided, a computational implementation would be highly desirable. Hopefully this initial sketch will lead to the development of such implementations.

3. RE-ENACTMENT, SIMULATORS AND SIMULATIONS

The central constructs in this approach to abstraction are simulators and simulations. Before they can be defined, it is first necessary to introduce the mechanism of modality-specific re-enactment.

(a) *Re-enactment in modality-specific systems*

The basic idea behind this mechanism is that association areas in the brain capture modality-specific states during perception and action, and then reinstate them later to represent knowledge. When a physical entity or event is perceived, it activates feature detectors in the relevant modality-specific areas. During visual processing of a car, for example, populations of neurons fire for edges, vertices and planar surfaces, whereas others fire for orientation, colour and movement. The total pattern of activation over this hierarchically organized distributed system represents the entity in vision (e.g. Zeki 1993; Palmer

1999). Similar distributions of activation on other modalities represent how the entity feels and sounds, and the actions performed on it. A related account can be provided for introspective states that arise during the event. Patterns of activation in the amygdala and orbito-frontal areas, for example, represent emotional and affective reactions.

Once a pattern becomes active in a feature area, conjunctive neurons in an association area store the pattern's features for later use. Damasio (1989) refers to these association areas as 'convergence zones', and assumes that they exist at multiple hierarchical levels, ranging from posterior to anterior in the brain. Simmons & Barsalou (2003) present a more developed account of the convergence zone architecture that explains lesion-based deficits of categorical knowledge.

The convergence zone architecture has the functional ability to re-enact sensory-motor and introspective states: once conjunctive neurons in a convergence zone capture a pattern of activation in a feature area, these neurons can later reinstate the pattern in the absence of bottom-up stimulation. During the recollection of a perceived object, for example, conjunctive neurons re-enact the sensory-motor and introspective states that were active while processing it originally. During the conceptualization of a category, conjunctive neurons similarly re-enact the modality-specific states characteristic of its members. No re-enactment is ever complete, and various biases may distort its reactivation. However, at least some semblance of the original state is partly reinstated.

This basic mechanism is widely viewed as underlying mental imagery (e.g. Kosslyn 1994; Jeannerod 1995; Farah 2000; Halpern 2001). The re-enactments it produces, however, are not necessarily conscious mental images. On the contrary, *unconscious* re-enactments may often underlie memory, conceptualization, comprehension and reasoning (Barsalou 1999). Although explicit attempts to construct mental imagery may create vivid re-enactments, other cognitive processes may often rely on less conscious re-enactments. In the account of abstraction developed here, the neural re-enactment of modality-specific states is the critical mechanism—*not* the re-enactment of conscious mental images.

(b) *Simulators*

According to Barsalou (1999), the neural re-enactment of modality-specific states underlies the conceptual system. In this theory, *simulators* and *simulations* constitute the central constructs. As different members of the same category are encountered, they generally activate similar neural patterns in feature areas (i.e. the statistically correlated features in categories; Rosch & Mervis (1975)). As a consequence of these shared feature activations, similar populations of conjunctive neurons tend to store these patterns in topographically related areas (Simmons & Barsalou 2003). As the category is learned, conjunctive neurons integrate sensory-motor and introspective features across its members, establishing a multi-modal representation of a category. Consider the category of *CARS*. Visual information about how cars look is integrated with auditory information about how they sound, olfactory information about how they smell, motor information about driving them, somatosensory information about feeling the ride in them, and emotional information associated with speed, dangerous situations, etc. The resulting representation is a distributed system throughout the brain's association and modality-specific areas that establishes knowledge about *CARS*. In Barsalou (1999), this distributed system is called a *simulator*.

(c) *Simulations*

A simulator is *not* a static representation of a category. Instead, it is a generator of representations. Specifically, a simulator re-enacts small subsets of its content as specific *simulations* on particular occasions to represent the respective category. The simulator's entire content is never activated all at once—only a small subset becomes active that is tailored to the constraints of the current situation (cf. Barsalou 1987, 1989, 1993). As Barsalou (2003) proposes, the active subset is configured to support the current course of situated action, providing goal-relevant inferences about objects, actions, mental states and the background setting. On one occasion, the *CAR* simulator might produce a simulation of travelling in car, whereas on others it might produce simulations of repairing a car, seeing a car park and so forth.

Simulations support a wide variety of functions in the cognitive system. For example, simulations produce inferences about category members that go beyond the information perceived for them. More generally, simulations constitute the representations that underlie memory, language and thought. See Barsalou (1999, 2003) for further detail.

(d) *Property simulators*

An infinitely large number of simulators can become established in the brain. Barsalou (1999) proposed that a simulator develops for any component of experience that attention selects repeatedly (also Mandler 1992). If attention focuses repeatedly on a particular component of experience across occasions, a simulator comes to represent it. As a result, simulators develop for various types of object, location, event, action, mental state and so forth. The flexibility of acquiring simulators is consistent with the argument of Schyns *et al.* (1998) that new features can be learned creatively. As these features become relevant for categorization, attention focuses on them, such that modality-specific information extracted from them becomes integrated into memory.

In the theory developed here, property simulators and relation simulators are central to the abstraction process. Each is addressed in turn. A property simulator arises from repeatedly processing a property of a category's members. Consider the property of *noses*. As attention focuses on the noses of a category's members (e.g. *HUMANS*), feature areas represent the relevant sensory-motor features in vision, and convergence zones capture these patterns. Later, in the absence of seeing a nose, the *nose* simulator can produce many different simulations of noses.

This account assumes that property simulators have the following four characteristics (for further detail, see Barsalou 2004). First, these simulators capture multi-modal information and produce multi-modal simulations. Simulations of *noses* contain not just visual information but also information from other relevant modalities, such as audition, smell and movement. Second, property simulations are typically not constructed in isolation. Rather than simulating a nose in isolation, it is typically simulated in the context of a background object or setting, such as a face. Third, these simulations typically do not produce global representations of a property that cover its form across all relevant categories. Rather than simulating a highly schematic nose, the *nose* simulator typically simulates specific noses, such as those for humans, dogs, fishes, and aeroplanes. Fourth, these specific simulations of a property are organized in a dominance order, such that some simulations are more likely to become active than others. For *noses*, simulations of human noses are most likely, with dog noses and aeroplane noses decreasing in accessibility. Solomon & Barsalou (2001) provide evidence for this account of property simulators. Related findings can also be found in Halff *et al.* (1976), Wisniewski (1998), Martin *et al.* (2000), Martin (2001), Martin & Chao (2001), Kan *et al.* (2003), Pecher *et al.* (2003*a*,*b*), Solomon & Barsalou (2003) and Wu & Barsalou (2003).

In this theory, property simulators develop for the wide variety of properties that people learn about categories. Thus, property simulators develop not only for visual properties (e.g. *noses*), but also for auditory properties (e.g. *barking*), motor properties (e.g. *pat*), touch properties (e.g. *soft*), smell properties (e.g. *sweet*), emotional categories (e.g. *happy*) and so forth.

(e) *Relation simulators*

As just demonstrated, a property simulator represents some aspect of a category's members. Analogously, a relation simulator represents *multiple* aspects of a category's members and their configuration (for further details, see Barsalou 2004). Consider the *above* relation. For the category of *FACES*, people not only acquire property knowledge about *noses* and *mouths*, they also acquire the knowledge that noses are *above* mouths. They also learn that *above* applies to many other property configurations, such as *roofs* being above *walls* in *HOUSES*, and *branches* being above *roots* in *TREES*. In all cases, two regions of the respective objects are relevant, where the focal region is higher vertically than the non-focal region.

As people repeatedly process these spatial relations, information about *above* accumulates in the *above* simulator. Information about the regions is extracted from experience, filtering out the respective entities they contain (e.g. the regions containing noses and mouths are extracted, filtering the details about noses and mouths). Once the simulator for *above* exists, it can produce many different simulations, each representing one particular configuration of regions for *above* (Talmy 1983; Langacker 1986; Herskovits 1997; Barsalou 1999). Many specific simulations are possible that vary in the vertical distance between the regions, the extent to which they are offset horizontally, their relative sizes, shapes and so forth. Although the details of these regions are filtered during the extraction of spatial information, different spatial patterns become stored with the respective categories. Thus the configuration of regions typical for *above* (*nose*, *mouth*) becomes associated with *FACES*, whereas the configurations typical for *above* (*roof*, *walls*) and *above* (*branches*, *roots*) become associated with *HOUSES* and *TREES*, respectively. As for property simulators, dominance orders of spatial configurations result, with some configurations being more accessible than others.

Increasing research supports the view that the meaning of a spatial preposition is a simulation of spatial regions. In research on attention, participants view a reference point, R, and assess whether another object, O, stands in some spatial relation to it (e.g. O is *above* R). O is then shown in many different positions around R, and the participant assesses the relation between them. Using typicality judgement and verification time as measures, researchers have shown that a prototypical configuration of spatial regions underlies the meaning of a spatial preposition (e.g. Hayward & Tarr 1995; Logan & Compton 1996; Carlson-Radvansky & Logan 1997). For example, the prototypical configuration for *above* occurs when the centre of O is aligned geometrically above the centre of R, not too far away. When participants hear 'above,' they construct a perceptual simulation of this prototypical configuration. When the display configuration subsequently matches this simulation, processing is optimal. As the configuration departs increasingly from the prototypical one, processing efficiency falls off in a graduated manner. The greater the departure from the prototypical configuration, the greater the transformation necessary to match it. Further research indicates that function modulates the particular spatial configuration simulated on a particular occasion (e.g. Coventry 1998; Carlson-Radvansky *et al.* 1999). As demonstrated for property simulators, a family of simulations—not just one simulation—underlies a spatial relation.

Besides spatial relations, a wide variety of other relations underlie human knowledge, including temporal, causal and intentional relations. Similar analyses can be applied to these other relations, namely, a simulator develops to produce a family of simulations for each relation (Barsalou 1999). As attention repeatedly focuses on the relevant configuration of regions across category members, a convergence zone captures the configuration and integrates it in a simulator. Temporal and causal relations develop from storing configurations of regions distributed over time, and intentional relations develop from storing configurations of regions that link mental states to the environment. Clearly, much further research is required to specify how these relations are learned and represented, but the basic idea is simulators develop to interpret the relevant configurations of spatio-temporal regions.

(f) *Holistic simulations*

Before returning to the issue of abstraction, one further construct must be established. So far, we have focused on how property and relation simulators develop to simulate the regions of entities and events. As will become clear shortly, entities and events must also be simulated, not just properties and relations within them. To verify that a *DOG* has a *nose* (in the absence of an actual dog), a dog must be simulated, not just a nose. The approach developed here is that the requisite conceptualizations of entities and events are holistic simulations built largely from global information extracted during pre-attentive processing.

Various types of global information could underlie these holistic simulations. For example, holistic representations could include blob-like representations of an entity's global shape, extracted by low-spatial frequency filters (e.g. Smith 1989; De Valois & De Valois 1988). Holistic representations could also include primary axes, parsed sub-regions and distributed configural features that capture direction and distance relations between sub-regions (e.g. Tanaka & Farah 1993). Thus, the basic idea is that when people simulate a *DOG* to verify that it has a *nose*, they simulate the dog using blob-like representations, structured by large configural features.

It is further assumed that these holistic simulations do not explicitly contain analytical properties of the sort described earlier for property and relation simulators. Holistic simulations do *not* explicitly represent properties at the conceptual level—they only contain perceptual information. For example, a holistic simulation of *AEROPLANE* might include perceptual information about wings but not explicitly represent the proposition *has* (*AEROPLANE, wings*). Instead, this proposition would only exist after deliberately binding the *wings* simulator to the holistic *AEROPLANE* simulation.

Nevertheless, property and relation simulators may implicitly influence the construction of holistic simulations without these relations becoming established explicitly. Specifically, property simulators that are highly associated with a category may implicitly influence a developing holistic simulation by articulating and enhancing regions that contain the respective properties. As a holistic simulation of *AEROPLANE* is constructed, a highly associated simulator for *wings* might influence the respective regions of the simulation. In general, however, the explicit representation of this property requires a more deliberate and conscious attempt to establish a relation between the two.

(g) *The DIPSS theory*

The previous sections laid the groundwork for the DIPSS theory of abstraction. There are no static summary representations in DIPSS, as in classic theories of abstraction. Nor are there summary representations that perfectly describe all of a category's members. Instead, the structural component of DIPSS is simply a loose collection of property and relation simulators. For example, as people learn about *BIRDS,* simulators for *wings*, *beak*, *feathers*, *nests* and *flies* develop to represent properties and relations important for this category.

The collection of property and relations simulators associated with a category is loose in the sense that it does not constitute a tight, structurally coherent theory. Rather than developing as an integrated formal system, these collections develop somewhat haphazardly, as various properties and relations become apparent for the category. Clearly, some simulators may develop together as one approach to analysing a category's exemplars, but they do not necessarily form an integrated formal system, such as a scientific theory. As will be demonstrated, these loose collections of property and relation simulators have the potential to explain the three properties of abstraction discussed earlier: type–token interpretation, structured representation and dynamic realization.

(h) *Type–token interpretation*

As we have seen, interpretation arises when concepts interpret components of experience. For example, when the concept for *CAR* becomes bound to a perceived visual entity, a type–token relation results that interprets the entity as a *CAR*. Once the type–token relation exists, inferences follow from the type to the token (e.g. the entity uses gasoline). Furthermore, this interpretation process is open ended, given that many different concepts can be applied to the same perceived entity.

Property simulators provide a natural account of type–token interpretation. Imagine that someone is examining a holistic simulation of a car (or a perceived one), scanning across it and describing its content. Over the course of this process, property simulators for *wheels*, *doors* and *windows* become bound to relevant regions of the simulated car via content addressable memory mechanisms.

Once a property simulator becomes bound to a region of a holistic simulation, an implicit type–token relation exists. The region of the holistic simulation is established as a token of the type that the property simulator represents. Mapping the *wheels* simulator into a region of a simulated car types the region as a *wheel*. The result is an implicit proposition that could be either true or false, and that carries inferences from the type to the token. For example, the *wheel* simulator might produce simulated inferences about the token rolling or going flat. Specifically, the *wheel* simulator might simulate a wheel rolling or going flat that goes beyond the information in the wheel region of the simulated car. Via the simulation mechanism, the standard categorical inferences associated with type–token propositions follow once simulators become bound to relevant regions of holistic simulations (Barsalou 1999).

In principle, a holistic simulation is subject to infinite property interpretations (as is a perceived category member). Because holistic simulations have a somewhat continuous quality, an infinite number of regions can be interpreted as properties. Furthermore, an infinite number of simulators could interpret a given region truly or falsely. This open-ended character is relevant to the later account of dynamic realization.

(i) *Structured representation*

In DIPSS, structured representation is simply a more complex form of type–token interpretation. By using relation simulators to interpret multiple regions of a holistic simulation simultaneously, structured representations result naturally. Consider how the following structured proposition about a *FACE* could be represented in perceptual symbol systems:

above(*upper-region* = *nose*(object-27),
    lower-region = *mouth*(object-41)).

The interpretative process begins by binding property simulators for *nose* and *mouth* to the respective regions of a face, thereby establishing two simple type–token relations. The relation simulator for *above* then becomes bound to the regions containing the nose and mouth regions, thereby interpreting both the regions and their contents as being in an *above* relation. In turn, this structured proposition could enter into a still more complex proposition that interprets the nose–mouth configuration as being below the eyes. As this example illustrates, when multiple simulators interpret a relational configuration of properties, a structured representation results (see Barsalou (1999) for further discussion). In this manner, perceptual symbol systems implement structured representations naturally and powerfully.

4. DYNAMIC REALIZATION

As we have seen, DIPSS assumes that people use loose collections of property and relation simulators to interpret category members. Although a relatively fixed set of property and relation simulators may exist for a person at a given point in time, the particular ones used across occasions vary considerably. Rather than being applied identically to different category members, these simulators are applied dynamically.

Consider the category of *CARS*. On one occasion, an assembly of property and relation simulators might be constructed to interpret how a car moves (e.g. using property simulators for *engine* and *wheels*). On another occasion, an assembly of simulators might be constructed to interpret the pleasure experienced while driving (e.g. using property simulators for *seat* and *stereo*). On yet another occasion, an assembly of simulators might be constructed to interpret how a car could prevent injury in a crash (e.g. using simulators for *seatbelt* and *airbag*). In each case, a different assembly of property and relation simulators interprets a category member.

The subset of property and relation simulators that interprets a category member on a given occasion can be viewed as an abstraction. Besides interpreting the member, this abstraction classifies the category member implicitly as something that the abstraction covers. Notably, these abstractions are not the classic sort of summary representation found in traditional theories. Once the category member drops from attention, the abstraction that interpreted it becomes largely irrelevant. The next time this member or another is processed, a different abstraction may be constructed dynamically to interpret it. Thus, abstractions are temporary online constructions, derived from a loose set of property and relation simulators used to interpret category members.

(a) *Interpretive attractors*

As we just saw, the abstractions constructed to interpret a category's members vary widely. Nevertheless, these abstractions are not constructed randomly from available property and relation simulators. Because of frequency and recency, some simulators may be more likely to be applied than others. Simulators applied frequently in the past will have an advantage, as will simulators applied recently. Furthermore, associations between simulators may produce correlations in the simulators assembled to interpret category members. Similarly, particular interpretive strategies may become associated with different contexts. As a result of such factors, interpretive attractors develop.

Across occasions, both statistical attractors and dynamic variability characterize the abstraction process. From this perspective, abstraction is more of a skill than a structure. As people learn about a category, they learn to interpret the properties and relations of its members, storing this knowledge in the respective simulators. As the skill for abstraction develops, a person can effectively process more regions of category members, and know the most appropriate regions to process in a particular context. The long-term outcome of this process is not a fixed summary representation, but a dynamic skill for interpreting category members effectively and efficiently.

(b) *Summary of the DIPSS approach to abstraction*

We began with three properties of abstraction: type–token interpretation, structured representation and dynamic realization. DIPSS naturally explains all three. Type–token interpretation results from applying property and relation simulators to the regions of perceived and simulated entities. Once these mappings exist, simulators produce inferences via the simulations they produce. Structured representation results from the simultaneous and integrated interpretation of a category member using multiple simulators. The classic processes of argument binding and recursion arise naturally in this process. Finally, dynamic realization results from the online application of a loose collection of property and relation simulators to a category's members. On a given occasion, a subset of simulators interprets a member, producing a temporary online abstraction. Across occasions, the abstractions constructed vary widely.

The diversity of the resulting abstractions explains various problems associated with classic theories. No one abstraction can be identified and motivated as *the* summary representation of a category, because an infinite number are possible. Furthermore, none of these abstractions needs to provide a complete account of the category. Instead, each abstraction interprets just those aspects of a category member that are currently relevant.

5. APPLICATIONS OF DIPSS TO ABSTRACTION PHENOMENA

Barsalou (2004) describes how DIPSS can be applied to various classes of abstraction phenomena. In categorization, DIPSS provides new ways of thinking about holistic versus analytical processing, dimension weights in exemplar models and the problem of descriptive inadequacy. On the topic of inference, DIPSS provides new ways of thinking about feature listing, conceptual instability, script tracks and verbal overshadowing. In the area of background knowledge, DIPSS provides new ways of thinking about intuitive theories, dimensions, multi-dimensional spaces and analogy. In learning, DIPSS provides new ways of thinking about shallow explanation, expertise and conceptual change. The brief summary here reviews only one example from each class of phenomena to illustrate how DIPSS explains abstraction phenomena (see Barsalou (2004), for descriptions of the others).

(a) *Descriptive inadequacy in categorization*

Philosophers often note the difficulty of specifying the properties that define a category (e.g. Wittgenstein 1953). Putnam (1973, 1975) argued that whatever description a person has for a category, it will never be sufficient to fix the category's reference (e.g. also Fodor 1998; Margolis & Laurence 1999). If the category's description turns out to be inadequate, the category's reference does not change, indicating that something besides the description establishes membership. If *WATER* turns out not to have the property H_2O but has some other property instead, the entities classified as *WATER* nevertheless remain constant. The property H_2O did not fix reference and was an inadequate description of category members.

DIPSS explains descriptive inadequacy as follows. Descriptions of a category are abstractions that arise from applying property and relation simulators in the available pool. Because this pool evolves haphazardly, and because descriptions are constructed dynamically, these descriptions never fully fix the category. Indeed, DIPSS embraces descriptive inadequacy. Abstraction is simply a skill that supports goal achievement in particular situations. It does not construct summary representations that fix category membership.

What does fix a category's reference as descriptions about it vary? As the lay understanding of *WATER* evolved with scientific theories, why did the reference of *WATER* remain basically the same? DIPSS explains this as the result of pre-attentive holistic representations of the sort described earlier for holistic simulations. Low-level sensory representations of *WATER* are likely to remain relatively constant as properties and relations about *WATER* change. As beliefs about *WATER* come and go, the perception of *WATER* remains relatively constant. Because these perceptions tend to be accurate in fixing category membership, holistic representations of them play the central role in everyday categorization, regardless of the analytical properties currently used for interpretation. These holistic representations could underlie the causal links that many philosophers propose are central to establishing reference.

Although analytical properties do not fix categorization, they may nevertheless influence it. In Biederman & Shiffrar (1987), participants learned analytical properties for chicken genitalia that facilitated their ability to categorize male versus female chicks. In Lin & Murphy (1997), participants learned functions for artificial objects that influenced their ability to categorize the objects visually. Such findings illustrate that holistic representations are not the sole determinants of categorization. Property and relation simulators also influence categorization, even when they do not completely fix it.

(b) *Feature listing as a form of category inference*

Researchers have often assumed that the feature-listing task provides a window on the underlying summary representation of a category. When participants generate a category's features, they presumably access a feature list, semantic network or schema for a category, and then read out the information verbally. For example, producing features for *CHAIRS* accesses a semantic network that specifies *seat*, *back* and *legs* as some of the underlying summary features.

From the perspective of DIPSS, no such underlying abstractions exist. Instead, participants construct a holistic simulation of the target category (e.g. a particular chair), and then interpret this simulation using property and relation simulators (e.g. property simulators for *seat*, *back* and *legs*). Instead of measuring a category's underlying summary representation, feature listing simply reflects one of many possible temporary abstractions that can be constructed online to interpret a particular member. Measuring these abstractions can be useful and informative (e.g. Wu & Barsalou 2003). They should not, however, be viewed as describing an underlying summary representation that covers the category or that fixes its reference. Instead, these feature listings are simply online inferences about a few properties of a category's members, based on the interpretation of one particular simulation. Because of the dynamical nature of feature listing, considerable variability arises both between and within individual people in the features they produce (e.g. Barsalou 1987, 1989, 1993).

(c) *Intuitive theories in background knowledge*

In a classic paper, Murphy & Medin (1985) proposed that intuitive theories provide background knowledge for categories. Problematically, however, little progress has been made in formulating these theories, and little agreement exists on the form they should take. DIPSS explains this quandary and provides a solution to it. Just as there is no single abstraction for a category, there is no single intuitive theory. From the perspective of DIPSS, a category's background knowledge is simply the loose collection of property and relation simulators used to interpret its members, together with the skill to apply them appropriately.

In some task contexts, this interpretive system may produce online abstractions along the lines of an intuitive theory. For example, property and relation simulators could be configured to explain how biological mechanisms maintain the life of an organism. On another occasion, a different abstraction might be constructed to explain an organism's reproductive origins. Across different occasions, different explanatory accounts of a category's members may be constructed (e.g. Gutheil *et al.* 1998). Thus, no single intuitive theory constitutes the background knowledge of a category. Instead, a loose collection of property and relation simulators produces theory-like abstractions (among others) dynamically to interpret category members.

(d) *Expertise in learning*

As learners become knowledgeable about a category, their stock of property and relation simulators for interpreting its members grows. Over time, this increasingly sophisticated interpretive system produces abstractions that are increasingly deep and useful (cf. Chi *et al.* 1981). Experts can interpret more critical regions in category members, and they can structure their interpretations in more sophisticated manners. As experts become adept at interpreting and organizing the regions of category members, they also become better categorizers, shifting their basic level down to the subordinate level (e.g. Johnson & Eilers 1998; Johnson & Mervis 1997, 1998; Gauthier *et al.* 2000).

Theories of expertise generally assume that expert performance results from the increased storage of exemplars, chunks or rules (e.g. Anderson 1987; Logan 1988; Newell 1990). In DIPSS, expert performance also results from the accumulation of property and relation simulators, and from increasing skill at applying them. The development of expertise can also be viewed as the accumulation of interpretive attractors for many different category members. Over time, all the relevant configurations for interpreting a category's members establish attractors in memory, leading to the relatively effortless performance that characterizes expert performance.

6. SIX SENSES OF ABSTRACTION REVISITED

One particular sense of abstraction—abstraction as summary representation—has been the focus of this article. The DIPSS account of this sense can now be brought to bear on the five other senses reviewed earlier.

(i) *Categorical knowledge*. According to this sense, knowledge of a category is abstracted from the experience of its members. In DIPSS, this amounts to establishing property and relation simulators that can interpret regions of perceived members and holistic simulations of them.

(ii) *The behavioural ability to generalize across category members*. According to DIPSS, when people behaviourally state a generic, such as 'cats have fur,' they have simulated a variety of *CAT* instances, used the *fur* simulator to interpret these simulations, and then used language to describe the temporary online abstraction.

(iii) *Summary representation*. Once a temporary abstraction is constructed for a category, a trace of it becomes established in memory, increasing the likelihood of constructing the abstraction later on another occasion. This abstraction, however, does not become part of a single summary representation for the category. It simply changes the dynamic qualities of the interpretive system, moving it towards an attractor. Nevertheless, the system remains dynamic, such that future abstractions vary widely, each tailored to the current situation and to the statistics of the interpretive system.

(iv) *Schematic representation*. Summary representations are schematic in the sense that they abstract the critical properties of category members and discard irrelevant ones. DIPSS accomplishes this 'sparseness of representation' in three ways. First, the property and relation simulators that develop for a category do not exhaust the simulators possible but only constitute a limited set. The resulting interpretive system is therefore schematic, representing just some aspects of category members. Second, the abstractions that property and relation simulators represent typically contain far less information than the sensory-motor perceptions that produced them. They are therefore schematic because they re-enact partial information and discard details. Third, property and relation simulators can produce idealized or caricatured simulations, thereby being schematic in the sense of producing prototypical representations. Such representations could result from the passive integration or averaging of the information in a simulator, with the most prototypical category information emerging as a dominant simulation (cf. the echo of Hintzman (1986)).

(v) *Flexible representation*. According to this sense of abstraction, summary representations can be applied flexibly to many different tasks. In DIPSS, this flexibility does not result from a single abstracted representation, but from a dynamic interpretive system. As expertise with a category develops, the set of property and relations simulators increases, as does the skill in applying them to members. As a result, the flexibility of interpretation increases, although attractors may produce entrenched ruts that work against flexibility to some extent.

(vi) *Abstract concepts*. According to this final sense of abstraction, some concepts become increasingly disengaged from physical entities and become increasingly associated with mental events (e.g. *truth*). In an analysis of abstract concepts, K. Wiemer-Hastings (unpublished data) found that many abstract concepts refer to properties and relations—not to complete objects and events—suggesting that abstract concepts belong to interpretive systems. This finding is consistent with the proposal of Barsalou (1999) that abstract concepts identify complex relational configurations of physical and mental states in background events.

In one sense of *truth,* for example, a complex relation exists between one person who makes a claim about the world and another person who assesses whether the claim is accurate. For *truth* to apply in a situation, a speaker must make a claim, a listener must represent it, the listener must compare this representation with the world and the representation must be accurate. When this complex relation is satisfied, *truth* is a valid interpretation of the speaker's claim. As this example illustrates, abstract concepts often capture complex configurations of physical and mental events. Similar to relation simulators, abstract concepts interpret multiple regions of events, and thus belong to the loose collections of simulators that constitute interpretive systems. Perhaps the distinguishing characteristic of abstract concepts is the complexity of the relational information they capture, together with their substantial inclusion of mental state information.

7. CONCLUSION

Interpretation and structured representations are two hallmarks of human cognition. The problem has been explaining these phenomena with mechanisms that exhibit dynamic realization instead of rigidity. Dynamic interpretation in perceptual symbol systems appears to offer a natural approach to unifying these three properties of abstractions. By applying loose collections of property and relation simulators to perceived and simulated category members, interpretation, structure and flexibility arise naturally in the abstraction process.

The author is grateful to Lorenza Saitta for the opportunity to publish this article and for assistance in preparing it. He is also grateful to Richard Patterson for helpful comments. This research was supported by National Science Foundation grants SBR-9905024 and BCS-0212134.

REFERENCES

Anderson, J. R. 1976 *Language, memory, and thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. 1987 Skill acquisition: compilation of weak-method problem situations. *Psychol. Rev.* **94**, 192–210.

Barsalou, L. W. 1985 Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J. Exp. Psychol. Learn. Mem. Cogn.* **11**, 629–654.

Barsalou, L. W. 1987 The instability of graded structure: implications for the nature of concepts. In *Concepts and conceptual development: ecological and intellectual factors in categorization* (ed. U. Neisser), pp. 101–140. Cambridge University Press.

Barsalou, L. W. 1989 Intraconcept similarity and its implications for interconcept similarity. In *Similarity and analogical reasoning* (ed. S. Vosniadou & A. Ortony), pp. 76–121. Cambridge University Press.

Barsalou, L. W. 1992 Frames, concepts, and conceptual fields. In *Frames, fields, and contrasts: new essays in semantic and lexical organization* (ed. E. Kittay & A. Lehrer), pp. 21–74. Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, L. W. 1993 Structure, flexibility, and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols. In *Theories of memory* (ed. A. C. Collins, S. E. Gathercole & M. A. Conway), pp. 29–101. London: Lawrence Erlbaum.

Barsalou, L. W. 1999 Perceptual symbol systems. *Behav. Brain Sci.* **22**, 577–660.

Barsalou, L. W. 2003 Situated simulation in the human conceptual system. Lang. Cogn. Process. (In the press.)

Barsalou, L. W. 2004 Abstraction as dynamic construal in perceptual symbol systems. In *Building object categories* (ed. L. Gershkoff-Stowe & D. Rakison). Carnegie Symposium Series. Mahwah, NJ: Lawrence Erlbaum. (In the press.)

Barsalou, L. W. & Hale, C. R. 1993 Components of conceptual representation: from feature lists to recursive frames. In *Categories and concepts: theoretical views and inductive data analysis* (ed. I. Van Mechelen, J. Hampton, R. Michalski & P. Theuns), pp. 97–144. San Diego, CA: Academic Press.

Biederman, I. 1987 Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* **94**, 115–147.

Biederman, I. & Shiffrar, M. M. 1987 Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. *J. Exp. Psychol. Learn. Mem. Cogn.* **13**, 640–645.

Bobrow, D. G. & Collins, A. M. (eds) 1975 *Representation and understanding: studies in cognitive science*. New York: Academic Press.

Carlson-Radvansky, L. A. & Logan, G. D. 1997 The influence of reference frame selection on spatial template construction. *J. Mem. Lang.* **37**, 411–437.

Carlson-Radvansky, L. A., Covey, E. S. & Lattanzi, K. M. 1999 'What' effects on 'where': functional influences on spatial relations. *Psychol. Sci.* **10**, 516–521.

Charniak, E. & McDermott, D. 1985 *Introduction to artificial intelligence*. Reading, MA: Addison-Wesley.

Chi, M. T. H., Feltovich, P. J. & Glaser, R. 1981 Categorization and representation of physics problems by experts and novices. *Cogn. Sci.* **5**, 121–152.

Church, A. 1956 *The problem of universals*. Notre Dame, IN: University of Notre Dame Press.

Collins, A. M. & Loftus, E. F. 1975 A spreading activation theory of semantic processing. *Psychol. Rev.* **82**, 407–428.

Coventry, K. R. 1998 Spatial prepositions, functional relations, and lexical specification. In *Representation and processing of spatial expressions* (ed. P. Oliver & K. P. Gapp), pp. 247–262. Mahwah, NJ: Lawrence Erlbaum.

Damasio, A. R. 1989 Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition* **33**, 25–62.

De Valois, R. & De Valois, K. 1988 *Spatial vision*. New York: Oxford University Press.

Dretske, F. 1995 *Naturalizing the mind*. Cambridge, MA: MIT Press.

Farah, M. J. 2000 The neural bases of mental imagery. In *The cognitive neurosciences*, 2nd edn (ed. M. S. Gazzaniga), pp. 965–974. Cambridge, MA: MIT Press.

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

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

Gauthier, I., Skudkarski, P., Gore, J. C. & Anderson, A. W. 2000 Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neurosci.* **3**, 191–197.

Gentner, D. & Markman, A. B. 1997 Structure mapping in analogy and similarity. *Am. Psychol.* **52**, 45–56.

Goldstone, R. L. & Medin, D. L. 1994 The course of comparison. *J. Exp. Psychol. Learn. Mem. Cogn.* **20**, 29–50.

Graesser, A. C., Singer, M. & Trabasso, T. 1994 Constructing inferences during narrative text comprehension. *Psychol. Rev.* **101**, 371–395.

Gutheil, G., Vera, A. & Keil, F. C. 1998 Do houseflies think? Patterns of induction and biological beliefs in development. *Cognition* **66**, 33–49.

Halff, H. M., Ortony, A. & Anderson, R. C. 1976 A context-sensitive representation of word meanings. *Mem. Cogn.* **4**, 378–383.

Halpern, A. R. 2001 Cerebral substrates of musical imagery. *Ann. NY Acad. Sci.* **930**, 179–192.

Hampton, J. A. 1997 Conceptual combination. In *Knowledge, concepts, and categories* (ed. K. Lamberts & D. Shanks), pp. 133–159. Cambridge, MA: MIT Press.

Hayward, W. G. & Tarr, M. J. 1995 Spatial language and spatial representation. *Cognition* **55**, 39–84.

Haugeland, J. 1991 Representational genera. In *Philosophy and connectionist theory* (ed. W. Ramsey, S. P. Stitch & D. E. Rumelhart), pp. 61–89. Hillsdale, NJ: Lawrence Erlbaum Associates.

Herskovits, A. 1997 Language, spatial cognition, and vision. In *Temporal and spatial reasoning* (ed. O. Stock), pp. 155–202. Norwell, MA: Kluwer.

Hintzman, D. L. 1986 'Schema abstraction' in a multiple-trace memory model. *Psychol. Rev.* **93**, 411–428.

Hochberg, J. 1998 Gestalt theory and its legacy: organization in eye and brain, in attention and mental representation. In *Perception and cognition at century's end: handbook of perception and cognition*, 2nd edn (ed. J. Hochberg), pp. 253–306. San Diego, CA: Academic Press.

Holte, R. C. & Choueiry, B. Y. 2003 Abstraction and reformulation in artificial intelligence. *Phil. Trans. R. Soc. Lond.* B **358**. (In this issue.) (DOI 10.1098/rstb.2003.1317.)

Holyoak, K. J. & Thagard, P. 1997 The analogical mind. *Am. Psychol.* **52**, 35–44.

Jeannerod, M. 1995 Mental imagery in the motor context. *Neuropsychologia* **33**, 1419–1432.

Johnson, K. E. & Eilers, A. T. 1998 Effects of knowledge and development on subordinate level categorization. *Cogn. Dev.* **13**, 515–545.

Johnson, K. E. & Mervis, C. B. 1997 Effects of varying levels of expertise on the basic level of categorization. *J. Exp. Psychol. Gen.* **126**, 248–277.

Johnson, K. E. & Mervis, C. B. 1998 Impact of intuitive theories on feature recruitment throughout the continuum of expertise. *Mem. Cogn.* **26**, 382–401.

Kan, I. P., Barsalou, L. W., Solomon, K. O., Minor, J. K. & Thompson-Schill, S. L. 2003 Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge. *Cogn. Neuropsychol.* (In the press.)

Kosslyn, S. M. 1994 *Image and brain*. Cambridge, MA: MIT Press.

Kintsch, W. & van Dijk, T. A. 1978 Toward a model of text comprehension and production. *Psychol. Rev.* **85**, 363–394.

Langacker, R. W. 1986 An introduction to cognitive grammar. *Cogn. Sci.* **10**, 1–40.

Lin, E. L. & Murphy, G. L. 1997 Effects of background knowledge on object categorization and part detection. *J. Exp. Psychol. Hum. Percept. Perform.* **23**, 1153–1169.

Logan, G. D. 1988 Toward an instance theory of automatization. *Psychol. Rev.* **95**, 492–527.

Logan, G. D. & Compton, B. J. 1996 Distance and distraction effects in the apprehension of spatial relations. *J. Exp. Psychol. Hum. Percept. Perform.* **22**, 159–172.

Mandler, J. M. 1992 How to build a baby: II. Conceptual primitives. *Psychol. Rev.* **99**, 587–604.

Margolis, E. & Laurence, S. 1999 Concepts and cognitive science. In *Concepts: core readings* (ed. E. Margolis & S. Laurence), pp. 3–81. Cambridge, MA: MIT Press.

Markman, A. B. & Gentner, D. 1997 The effects of alignability on memory. *Psychol. Sci.* **8**, 363–367.

Martin, A. 2001 Functional neuroimaging of semantic memory. In *Handbook of functional neuroimaging of cognition* (ed. R. Cabeza & A. Kingstone), pp. 153–186. Cambridge, MA: MIT Press.

Martin, A. & Chao, L. 2001 Semantic memory and the brain: structure and process. *Curr. Opin. Neurobiol.* **11**, 194–201.

Martin, A., Ungerleider, L. G. & Haxby, J. V. 2000 Category-specificity and the brain: the sensory-motor model of semantic representations of objects. In *The new cognitive neurosciences*, 2nd edn (ed. M. S. Gazzaniga), pp. 1023–1036. Cambridge, MA: MIT Press.

Murphy, G. L. & Medin, D. L. 1985 The role of theories in conceptual coherence. *Psychol. Rev.* **92**, 289–316.

Newell, A. 1990 *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Newell, A. & Simon, H. A. 1972 *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Paivio, A. 1986 *Mental representations: a dual coding approach*. New York: Oxford University Press.

Palmer, S. E. 1999 *Vision science: photons to phenomenology*. Cambridge, MA: MIT Press.

Palmeri, T. J. & Nosofsky, R. M. 2001 Central tendencies, extreme points, and prototype enhancement effects in ill-defined perceptual categorization. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* **54**, 197–235.

Pecher, D., Zeelenberg, R. & Barsalou, L. W. 2003*a* Verifying properties from different modalities for concepts produces switching costs. *Psychol. Sci.* **14**, 119–124.

Pecher, D., Zeelenberg, R. & Barsalou, L. W. 2003*b* Sensorimotor simulations underlie conceptual representations: modality-specific effects of prior activation. *Psychonom. Bull. Rev.* (In the press.)

Posner, M. I. & Keele, S. W. 1968 On the genesis of abstract ideas. *J. Exp. Psychol.* **77**, 353–363.

Putnam, H. 1973 Meaning and reference. *J. Phil.* **70**, 699–711.

Putnam, H. 1975 The meaning of 'meaning'. In *Mind, language, and reality: philosophical papers*, vol. 2 (ed. H. Putnam), pp. 215–271. New York: Cambridge University Press.

Pylyshyn, Z. W. 1973 What the mind's eye tells the mind's brain: a critique of mental imagery. *Psychol. Bull.* **80**, 1–24.

Rosch, E. & Mervis, C. B. 1975 Family resemblances: studies in the internal structure of categories. *Cogn. Psychol.* **7**, 573–605.

Rhodes, G., Brennan, S. & Carey, S. 1987 Identification and ratings of caricatures: implications for mental representation of faces. *Cogn. Psychol.* **19**, 473–497.

Rips, L. J. 1995 The current status of research on concept combination. *Mind Lang.* **10**, 72–104.

Schank, R. C. & Abelson, R. P. 1977 *Scripts, plans, goals, and understanding: an inquiry into human knowledge structures*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. C. & Colby, K. M. (eds) 1973 *Computer models of thought and language*. San Francisco, CA: Freeman.

Schyns, P. G., Goldstone, R. L. & Thibaut, J. P. 1998 The development of features in object concepts. *Behav. Brain Sci.* **21**, 1–54.

Simmons, K. & Barsalou, L. W. 2003 The similarity-in-topography principle: reconciling theories of conceptual deficits. *Cogn. Neuropsychol.* (In the press.)

Smith, L. B. 1989 A model of perceptual classification in children and adults. *Psychol. Rev.* **96**, 125–144.

Solomon, K. O. & Barsalou, L. W. 2001 Representing properties locally. *Cogn. Psychol.* **43**, 129–169.

Solomon, K. O. & Barsalou, L. W. 2003 Perceptual simulation in property verification. (Submitted.)

Talmy, L. 1983 How language structures space. In *Spatial orientation: theory, research, and application* (ed. H. Pick & L. Acredelo), pp. 225–282. New York: Plenum Press.

Tanaka, J. W. & Farah, M. J. 1993 Parts and wholes in face recognition. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* **42**, 225–245.

Wiemer-Hastings, K., Krug, J. & Xu, X. 2001 Imagery, context availability, contextual constraint, and abstractness. In *Proc. 23rd Ann. Conf. Cogn. Sci. Soc.*, pp. 1134–1139. Mahwah, NJ: Lawrence Erlbaum.

Winograd, T. 1972 *Understanding natural language*. San Diego, CA: Academic Press.

Winograd, T. 1975 Frame representations and the declarative-procedural controversy. In *Representation and understanding: studies in cognitive science* (ed. D. G. Bobrow & A. M. Collins). New York: Academic Press.

Wisniewski, E. J. 1997 When concepts combine. *Psychonom. Bull. Rev.* **4**, 167–183.

Wisniewski, E. J. 1998 Property instantiation in conceptual combination. *Mem. Cogn.* **26**, 1330–1347.

Wittgenstein, L. 1953 *Philosophical investigations*. New York: Macmillan [Translation by G. E. M. Anscombe.]

Wu, L., & Barsalou, L. W. 2003 Perceptual simulation in property generation. (Submitted.)

Zeki, S. 1993 *A vision of the brain*. Oxford: Blackwell Scientific.

Zucker, J.-D. 2003 A grounded theory of abstraction in artificial intelligence. *Phil. Trans. R. Soc. Lond.* B **358**. (In this issue.) (DOI 10.1098/rstb.2003.1308.)

GLOSSARY

DIPSS: dynamic interpretation in perceptual symbol systems