Chapter 28

SITUATED CONCEPTUALIZATION

LAWRENCE W. BARSALOU
Emory University

Contents

Handbook of Categorization in Cognitive Science, Edited by Henri Cohen and Claire Lefebvre

Abstract

Two themes about the conceptual system are developed: (1) modal simulations underlie conceptual processing; (2) conceptual representations are situated. The construct of situated conceptualization–a multimodal simulation that supports one specific course of situated action with a particular category instance–integrates these themes. A given concept produces many different situated conceptualizations, each tailored to different instances in different settings. A situated conceptualization creates the experience of "being there" with a category instance in a setting via integrated simulations of objects, settings, actions, and introspections. On recognizing a familiar type of instance, an entrenched situated conceptualization associated with it becomes active, which provides relevant inferences via pattern completion. Supporting empirical evidence from cognitive psychology, social psychology, and cognitive neuroscience is reviewed.

1. Introduction

1.1. Conceptual systems

The conceptual system is a system distributed throughout the brain that represents knowledge about the world. It is *not* a collection of holistic images like those in a camera, video recorder, or audio recorder. Instead, it is a collection of category representations, with each category corresponding to a *component* of experience–not to an entire holistic experience. Across a person's life span, category knowledge about such components develops for objects, locations, times, events, introspective states, relations, roles, properties, etc.

The conceptual system provides representational support across the spectrum of cognitive activities. In online processing, as people pursue goals in the environment, the conceptual system contributes in several important ways. First, it helps construct perceptions through figure-ground segregation, anticipation, and filling in. Second, it predicts entities and events likely to be perceived, speeding their processing. Third, it supports categorization, assigning perceived entities and events to categories. Fourth, it provides inferences following categorization that constitute expertise about the world. Rather than starting from scratch when interacting with an entity or event, agents benefit from knowledge of previous category members.

The conceptual system is also central to offline processing when people represent non-present entities and events in memory, language, and thought. In memory, the conceptual system provides elaboration at encoding, organizational structure in storage, and reconstructive inference during retrieval. In language, the conceptual system contributes to the meanings of words, phrases, sentences, and texts, and to the inferences that go beyond them. In thought, the conceptual system provides representations of the objects and events that are the objects of reasoning, decision making, and problem solving.

1.2. Semantic memory

Since the cognitive revolution, theorists have proposed many accounts of the conceptual system. The dominant theory, however, has been the semantic memory view, which arises from a proposed distinction between semantic and episodic memory [Tulving (1972)]. Specific models that instantiate this view include network models [e.g., Collins and Quillian (1969), Collins and Loftus (1975)] and feature set models [e.g., Rosch and Mervis (1975), Hampton (1979)]. For a review of semantic memory models, see Smith (1978). This approach to thinking about the conceptual system remains dominant. Researchers throughout the cognitive sciences continue to adopt various forms of semantic memory models in their working accounts of the cognitive system.

Four assumptions underlie the semantic memory view. First, semantic memory is viewed as a modular system, that is, as being autonomous relative to the episodic memory system and to the systems for perception, action, emotion, and motivation. From this theoretical perspective, the conceptual system does not share representation and processing mechanisms with these other brain systems, but is an independent module that operates according to different principles.

Second, and relatedly, semantic memory representations are assumed to be amodal, differing significantly from representations in modality-specific systems. Specifically, semantic memory representations are assumed to be *redescriptions* or *transductions* of modality-specific representations into a new representation language that does not have modality-specific qualities. Instead, these representations consist of arbitrary amodal symbols that stand for modality-specific states and for the entities in the world that these states represent.

Third, semantic memory representations are viewed as decontextualized. In the typical theory, the representation of a category is a prototype or definition that distills relatively invariant properties across exemplars. Lost in the distillation are idiosyncratic properties of exemplars and background situations. Thus the representation of [CAT] might be a decontextualized prototype that includes CLAWS, WHISKERS, and TAIL, with idiosyncratic properties and background situations filtered out[1]. As a result, category representations in the semantic memory view have the flavor of encyclopedia descriptions in a database of categorical knowledge about the world.

Fourth, semantic memory representations are typically viewed as being relatively stable. For a given category, different people share roughly the same knowledge, and the same person uses the same knowledge on different occasions.

2. Grounding the conceptual system in the modalities

A diametrically opposed way of thinking about the conceptual system is developed here. The first section of this chapter presents the theoretical assumptions of this approach; the second section presents empirical support for it.

In the theoretical section, the first of three subsections introduces the constructs of *reenactment*, *simulator*, and *simulation*. Rather than being modular, the conceptual system shares fundamental mechanisms with modality-specific systems. As a result, conceptual representations are modal, not amodal.

The second subsection introduces the construct of *situated conceptualization*, which grounds conceptual processing in situated action. Rather than being decontextualized and stable, conceptual representations are contextualized dynamically to support diverse courses of goal pursuit.

The third subsection illustrates how situated conceptualizations support conceptual inferences via pattern completion. When one part of a situated conceptualization is perceived, the remainder of the conceptualization becomes active, constituting inferences about the current situation.

[1] Following the conventions used throughout the Handbook, uppercase will be used to indicate conceptual representations, and italics will be used to indicate linguistic forms. Within conceptual representations, uppercase in brackets will indicate categories, whereas uppercase font with no brackets will indicate properties of categories. Thus, [CATS] indicates a category, whereas CLAWS indicates a property, with *cats* and *claws* indicating the respective linguistic forms.

2.1. Modal reenactments of perception, action, and introspection

The modal reenactment of perceptual, motor, and introspective states constitutes the central mechanism in this approach, not amodal redescriptions of these states [e.g., Damasio (1989), Barsalou (1999b, 2003a), Simmons and Barsalou (2003)]. The reenactment process underlying knowledge is assumed to be *approximately* the same as the reenactment process underlying mental imagery [e.g., Barsalou (1982), Finke (1989), Kosslyn (1994), Zatorre(1996), Farah (2000), Grezes and Decety et al (2001)]. The reenactment process has two phases: (1) the storage of modality-specific states, and (2) the partial reenactment of these states. Each phase is addressed in turn.

2.1.1. Storage of modality-specific states that arise in feature systems

When a physical entity is experienced, it activates feature detectors in the relevant brain systems. During visual processing of a cat, for example, neurons fire for edges and planar surfaces, whereas others fire for color, configural properties, and movement. The overall pattern of activation across this hierarchically organized distributed system represents the entity in vision [e.g., Zeki (1993), Palmer (1999)]. Analogous patterns of activation in other sensory modalities represent how the cat might sound and feel. Activations in the motor system represent actions on the cat. Similar mechanisms underlie the introspective states that arise while interacting with an entity. For example, activations in the amygdala and orbitofrontal areas might represent emotional reactions to the cat.

When a pattern becomes active in a feature system, conjunctive neurons in association areas capture the pattern for later cognitive use. A population of conjunctive neurons codes the pattern, with each individual neuron participating in the coding of many different patterns (i.e., coarse coding). Damasio (1989) calls these association areas *convergence zones*, and proposes that they exist at multiple hierarchical levels in the brain [also see Simmons and Barsalou (2003)]. Locally, convergence zones near a modality capture activation patterns within it. Association areas near the visual system capture patterns there, whereas association areas near the auditory system capture patterns there. Downstream in more anterior regions, higher association areas in the temporal, parietal, and frontal lobes integrate activation across modalities.

2.1.2. Reenactments of modality-specific states

The convergence-zone architecture has the functional ability to produce modality-specific reenactments. Once a set of conjunctive neurons captures a feature pattern, this set can later activate the pattern in the absence of bottom-up stimulation. When retrieving a memory of a cat, conjunctive neurons partially reactivate the visual state active during its earlier perception. Similarly, when retrieving an action performed on the cat, conjunctive neurons partially reactivate the motor state that produced it. A reenactment never constitutes a complete reinstatement of the original modality-specific state.

Furthermore, bias may often distort it. Thus, a reenactment is always partial and potentially inaccurate. Nevertheless, some semblance of the original state is reactivated–not an amodal redescription.

The reenactment process is not necessarily conscious. Although conscious reenactment is viewed widely as the process that underlies mental imagery, reenactments need not always reach awareness. Unconscious reenactments may often underlie memory, conceptualization, comprehension, and reasoning [Barsalou (1999b)]. Although explicit attempts to construct mental imagery may create vivid reenactments, many other cognitive processes may rely on less conscious reenactments, or reenactments that are largely unconscious.

2.2. Simulators and simulations

Barsalou (1999b, 2003a) developed a theory of the conceptual system based on the neural reenactment of modality-specific states. According to this view, a fully functional conceptual system can be built on reenactment mechanisms. Using these mechanisms, it is possible to implement the type-token distinction, categorical inference, productivity, propositions, and abstract concepts. Contrary to previous arguments, amodal symbols are not the only possible way to implement these classical functions of a conceptual system.

Simulators and *simulations* constitute the two central constructs of this theory. Simulators integrate information across a category's instances, whereas simulations are specific conceptualizations of the category. Each is addressed in turn.

2.2.1. Simulators

Categories tend to have statistically correlated properties [e.g., McRae, de Sa and Seidenberg (1997)]. Thus, encountering different instances of the same category should tend to activate similar neural patterns in feature systems [e.g., Farah and McClelland (1991), Cree and McRae (2003)]. Furthermore, similar populations of conjunctive neurons in convergence zones – tuned to these particular conjunctions of features – should tend to capture these similar patterns [Damasio (1989), Simmons and Barsalou (2003)]. Across experiences of instances and settings, this population of conjunctive neurons integrates modality-specific properties, establishing a multimodal representation of the category. Barsalou (1999b) refers to these distributed systems as *simulators*. Conceptually, a simulator functions as a type. It integrates the multimodal content of a category across instances, and provides the ability to interpret later individuals as tokens of the type [Barsalou (2003a)].

Consider the simulator for the category of [CATS]. Across learning, visual information about how cats look becomes integrated in the simulator, along with auditory information about how they sound, somatosensory information about how they feel, motor programs for interacting with them, emotional responses to experiencing them, and so forth. The result is a distributed system throughout the brain's feature and association areas that accumulates conceptual content for the category.

2.2.2. Simulations

Once a simulator becomes established for a category, it reenacts small subsets of its content as specific *simulations*. All of the content in a simulator never becomes active simultaneously. Instead, only a small subset becomes active to represent the category on a particular occasion [e.g., Barsalou (1987, 1989, 1993)]. For example, the [CAT] simulator might simulate a sleeping kitten on one occasion, whereas on other occasions it might simulate a hissing tom cat or a purring house cat. Because all the experienced content for cats resides implicitly in the [CATS] simulator, many different subsets can be reenacted on different occasions.

Simulations serve a wide variety of cognitive functions. As Barsalou (1999b, 2003a) illustrates, simulations can represent a category's instances in their absence during memory, language, and thought. Simulations can be used to draw inferences about a category's perceived instances using the pattern completion described later. They can be combined productively to produce infinite conceptual combinations. They also can represent the propositions that underlie type-token predication and recursion.

Simulations can also be used to represent novel category instances not already stored in a simulator. Instances stored on previous occasions may merge together at retrieval, thereby producing reconstructive and averaging effects. Remembering a cat seen once, for example, may be distorted toward a similar cat seen many times. Furthermore, intentional attempts to combine simulations of conceptual components can produce simulations never experienced. For example, people can simulate a cat and then systematically vary simulations of its color and patterning to represent a wide variety of novel instances.

2.2.3. Sources of simulators

In principle, an infinite number of simulators can develop in memory for all forms of knowledge, including objects, properties, settings, events, actions, introspections, and so forth. Specifically, a simulator develops for any component of experience that attention selects repeatedly [Barsalou (1999b, 2003a)]. When attention focuses repeatedly on a type of object in experience, such as for [CATS], a simulator develops for it. Analogously, if attention focuses on a type of action ([BRUSHING]) or on a type of introspection ([HAPPINESS]), simulators develop to represent them as well. Such flexibility is consistent with the Schyns, Goldstone and Thibant (1998) proposal that the cognitive system acquires new properties as they become relevant for categorization. Because selective attention is flexible and open-ended, a simulator develops for any component of experience selected repeatedly.

A key issue concerns why attention focuses on some components but not on others, such that simulators develop for those components. Many factors influence this process, including genetics, language development, culture, and goal achievement. A further account of these mechanisms lies beyond the scope of this chapter. Notably, though, this is the classic problem of what constrains knowledge (e.g., Murphy and Medin (1985)]. Any theory – not just this one – must resolve it.

Another key issue concerns how simulators for abstract concepts are represented. Barsalou (1999b) proposed that simulators for abstract concepts generally capture complex multimodal simulations of temporally extended situations, with simulated introspective states being central. Relative to concrete concepts, abstract concepts tend to contain more situational and introspective information than do concrete concepts. One sense of [TRUTH], for example, begins with a speaker making a claim about a situation, such as "It's sunny outside." A listener then represents the claim, compares it to the actual situation, and decides if the claim interprets the situation accurately. This sense of [TRUTH] can be represented as a simulation of the situation, including the relevant introspective states (e.g., representing, comparing, deciding). Many abstract concepts, such as [FREEDOM] and [INVENT], can similarly be viewed as complex simulations of situations, with simulated introspective states being central. Wiemer-Hastings, King and Xu (2001) and Barsalou and Wiemer-Hastings (2005) offer preliminary support for this proposal.

2.3. Situated conceptualizations

Barsalou (2003b) contrasts two ways of thinking about concepts [also see Barsalou (1999a)]. On the one hand, semantic memory theories implicitly view concepts as *detached databases*. When a category is learned, its properties and exemplars are integrated into a general description that is relatively detached from the goals of specific agents. On different occasions, a person uses the same general description to represent the category. Alternatively, a concept can be viewed as an *agent-dependent instruction manual* that delivers specialized packages of inferences to guide an agent's interactions with particular category members in specific situations. Across different situations, a concept delivers different packages of inferences, each tailored to current goals and constraints. Because a single general description would be too vague to support all the relevant inferences in a particular situation, more specialized representations are constructed instead.

Barsalou (2003b) referred to one particular package of situation-specific inferences as a *situated conceptualization*. Consider the concept of [CAT]. According to traditional views, [CAT] is represented as a detached collection of amodal facts that becomes active as a whole every time the category is processed. Alternatively, a simulator for [CAT] produces many different situated conceptualizations, each tailored to helping an agent interact with cats in a different context – no general description of the category exists. For example, one situated conceptualization for [CAT] might support interacting with a playful kitten, whereas others might support interacting with a mean tom cat, or with a purring house cat. In this view, the concept for [CAT] is not a detached global description of the category. Instead, the concept is the skill or ability to produce a wide variety of situated conceptualizations that support goal achievement in specific contexts.

2.3.1. Multimodal simulations implement situated conceptualizations

Barsalou (2003b) further proposed that a complex simulation becomes active across modalities to implement a situated conceptualization. Consider a situated conceptualization for interacting with a purring house cat. This conceptualization is likely to simulate how the cat might appear perceptually. When cats are purring, their bodies take particular shapes, they execute certain actions, and they make distinctive sounds. All these perceptual aspects can be represented as modal simulations in the situated conceptualization. Rather than amodal redescriptions representing these perceptions, simulations represent them in the relevant modality-specific systems.

A situated conceptualization about a purring house cat is likely to simulate actions that the agent could take in the situation, such as scratching the cat. Modal simulations can also represent these aspects of a situated conceptualization via simulations of the actions themselves, not by amodal redescriptions.

A situated conceptualization about a purring house cat is likely to include simulations of introspective states. Because people experience particular introspections around purring house cats, the respective situated conceptualizations include simulations of emotions, evaluations, motivations, cognitive operations, etc.

Finally, a situated conceptualization for a purring house cat simulates a setting where the event could take place – the event is not simulated in a vacuum. Thus, an interaction with a purring house cat might be simulated in a living room, bedroom, yard, etc. Again such knowledge is represented as simulations, this time as reenactments of particular settings.

In summary, a situated conceptualization typically simulates four basic types of components: (1) perceptions of relevant people and objects, (2) an agent's actions and other bodily states, (3) introspective states, such as emotions and cognitive operations, and (4) likely settings. Putting all these together, a situated conceptualization is a multimodal simulation of a multicomponent situation, with each modality-specific component simulated in the respective brain area.

It is important to note that a situated conceptualization consists of simulations from many different simulators. A situated conceptualization for a purring house cat is likely to include simulations from simulators for animals, people, objects, actions, introspections, and settings. Thus, a single simulator alone does not produce a situated conceptualization. Instead, many simulators contribute to the collection of components that a situated conceptualization contains.

It is also important to note that situated conceptualizations place the conceptualizer directly in the respective situations, creating the experience of "being there" [Barsalou (2002)]. By reenacting an agent's actions and introspective states, these complex simulations create the experience of the conceptualizer being in the situation – the situation is not represented as detached and separate from the conceptualizer.

2.3.2. Entrenched situated conceptualizations

Across their life spans, people experience many situations repeatedly in their interactions with people, artifacts, social institutions, etc. As a result, knowledge about these repeated situations becomes entrenched in memory, thereby supporting skilled performance in them. Entrenched knowledge can also guide interactions in novel situations that are similar to these familiar situations [e.g., Andersen and Chen (2002)]. Even though entrenched knowledge may not always provide a perfect fit, it may often fit well enough to provide useful inferences.

We assume that situated conceptualizations represent people's entrenched knowledge of these repeated situations. When a situation is experienced repeatedly, multimodal knowledge accrues in the respective simulators for the relevant people, objects, actions, introspections, and settings. The conceptualization's components become entrenched in the respective simulators, as do associations between these components. Over time, the situated conceptualization becomes so well established that it comes to mind automatically and immediately as a unit when the situation arises. After petting a purring house cat on many occasions, for example, the situated conceptualization for this situation becomes entrenched in memory, such that minimal cuing activates it on subsequent occasions.

2.4. Inference via pattern completion

Once situated conceptualizations become entrenched in memory, they play important roles throughout cognition. In perception, they support the processing of familiar scenes [e.g., Biederman (1981)]. In memory, they support reconstructive retrieval [e.g., Brewer and Treyens (1981)]. In language, they produce situation models and diverse forms of inference [e.g., Zwaan and Radvansky (1998)]. In reasoning, they provide content that facilitates deduction [e.g., Johnson-Laird (1983)]. In social cognition, they provide rich inferences about myriad aspects of interpersonal interaction [e.g., Barsalou et al. (2003)].

2.4.1. Pattern completion with entrenched situated conceptualizations

Much of the processing support that entrenched situated conceptualizations provides appears to result from a pattern–completion inference process. On entering a familiar situation and recognizing it, an entrenched situated conceptualization that represents the situation becomes active. Typically, not all of the situation is perceived initially. A relevant person, setting, or event may be perceived, which then suggests that a particular situation is about to unfold. It is in the agent's interests to anticipate what will happen next, so that optimal actions can be executed. The agent must draw inferences that go beyond the information given [e.g., Bruner (1957)].

The situated conceptualization that becomes active constitutes a rich source of inference. The conceptualization is essentially a pattern, namely, a complex configuration of multimodal components that represent the situation. When a component of this pattern matched the situation, the larger pattern became active in memory. The remaining pattern components–not yet observed–constitute inferences, that is, educated guesses about what might occur next. Because the remaining components co-occurred frequently with the perceived components in previous situations, inferring the remaining components is justified. When a partially viewed situation activates a situated conceptualization, the conceptualization completes the pattern that the situation suggests. To the extent that a situated conceptualization is entrenched in memory, this process is likely to occur relatively automatically.

Consider the example of seeing a particular cat. Imagine that her face, color, and bodily mannerisms initially match modality-specific simulations in one or more situated conceptualizations that have become entrenched in memory for [CATS]. Once one conceptualization wins the activation process, it provides inferences via pattern completion, such as actions that the cat is likely to take, actions that the perceiver typically takes, mental states that are likely to result, and so forth. The unfolding of such inferences–realized as simulations–produces inferential prediction.

2.4.2. The statistical character of inference

Everything about the production of inferences via pattern completion has a statistical character [e.g., Barsalou (1987, 1989, 1993), Smith and Samuelson (1997)]. Each simulator that contributes to a situated conceptualization is a dynamical system capable of producing infinite simulations [Barsalou (1999b, 2003a,b)]. In a particular situation, each simulation constructed reflects the current state of the simulator, its current inputs, and its past history. An entrenched situated conceptualization is essentially an attractor, namely, an associated collection of simulations that is easy to settle on, because the associations linking them have become strong through frequent use. Infinitely many states near the attractor, however, offer different versions of the same conceptualization, each representing a different adaptation to the situation. Thus, the entrenched conceptualization for interacting with a purring house cat is not just one complex simulation but the ability to produce many related simulations. When encountering the same type of situation on different occasions, the situated conceptualizations that guide an agent vary dynamically, depending on all relevant factors that influence the contributing simulators.

As a consequence, the inferences that arise via pattern completion vary as well. As the conceptualizations that represent a situation vary across occasions, the completions that follow also vary. Somewhat different inferences result from completing somewhat different patterns.

3. Empirical evidence

This section reviews evidence for the two theses of the previous section. On the one hand, accumulating findings implicate modality-specific mechanisms in conceptual processing. On the other, a variety of additional findings suggest that conceptual representations are situated.

3.1. Behavioral evidence for a modal nonmodular conceptual system

Defining the construct of a *modality-specific state* is useful for assessing whether modal systems underlie conceptual processing. One way to think about such states is as patterns of neural activation. On seeing a cat, its visual representation in the brain includes patterns of neural activation along the ventral and dorsal streams. Thinking about modality-specific states this way is well established and widely accepted [e.g., Zeki (1993), Palmer (1999)]. Alternatively, modality-specific states can be viewed as conscious mental images. Problematically, though, what becomes conscious is a relatively small subset of the unconscious processing occurring neurally. For this reason, the remainder of this chapter focuses on the neural representation of modality-specific states.

3.1.1. Predictions for modular amodal vs. nonmodular modal theories

Viewing modality-specific states as active neural patterns provides a means of distinguishing theses two approaches. According to modular amodal views, a conceptual representation is not a neural pattern in modality-specific systems. The neural patterns that represent an entity during its perception have nothing to do with its conceptual representation. Instead, neural patterns in another brain system represent the object conceptually, using a different format than those in modality-specific systems. Furthermore, amodal representations use the same general format to represent conceptual information about properties from different modalities. Thus, modular amodal views do not predict a *priori* that modality-specific systems should become active during conceptual processing, nor that different patterns of modality-specific processing should arise for categories having different modality-specific content.

Nonmodular modal views make the opposite predictions. During the conceptual representation of a category, the neural systems that process actual interactions with its instances should also become active as if a category member were present (although not identically). On conceptualizing [CATS], for example, the visual system might become partially active as if a cat were present. Similarly, the auditory system might reenact states associated with hearing a cat, the motor system might reenact states associated with petting a cat, and the limbic system might reenact emotional states associated with enjoying the experience of a cat.

Thus, modular amodal and nonmodular modal theories make different predictions about the roles of modality-specific systems in conceptual processing. Because of these clear differences, it seems that much research in the literature would have addressed which account is correct. Surprisingly, though, little research has addressed this issue. Instead, theoretical considerations have primarily been responsible for the widespread acceptance of modular amodal views throughout the cognitive science community. As Barsalou (1999b) conjectured, the ascendance of the modular amodal approach reflected the development of logic, statistics, and computer science in the twentieth century, and the subsequent incorporation of these developments into the cognitive revolution. Because amodal representation languages have much expressive power, because they can be formalized, and because they can be implemented in computer hardware, they captured the imagination of the cognitive science community, took over theoretical thinking, and became widely practiced.

Nevertheless, a strong empirical case should exist for such a central assumption, even if it is useful theoretically. As the next two subsections illustrate, accumulating findings question this assumption, first from behavioral psychology, and second, from cognitive neuroscience.

3.1.2. Assessing the presence of modality-specific effects in conceptual processing

The behavioral experiments reviewed here used laboratory tasks that are widely assumed to activate and utilize category knowledge. For example, the research in my laboratory has often studied the property generation and property verification tasks. During property generation, a participant hears the word for a category (e.g., *cat*), and then states characteristic properties of the underlying concept out loud (e.g., *claws, whiskers, tail, you scratch it*). During property verification, a participant reads the word for a category on a computer (e.g., *cat*), and then verifies whether a subsequently presented property is true or false for that category (e.g., *claws* vs. *wings*). Whereas property generation is an active, production-oriented task extended over time, property verification is a more passive, recognition-oriented task executed under time pressure.

Most accounts of these tasks assume that participants use amodal representations to perform them [e.g., Kosslyn (1976), Smith (1978)]. When producing or verifying properties, participants access semantic networks, feature lists, frames, etc., to produce the required information. We hypothesized instead that participants simulate a category member to represent a category, and then consult their simulations to produce the requested information. During property generation, participants scan across their simulations, and produce words for properties perceived in them [see Barsalou (2003a), for an account of how properties in the regions of a simulation are perceived]. During property verification, participants evaluate whether test properties can be perceived in these simulations.

Most importantly, these experiments employ the following logic to test the simulation hypothesis: If conceptual processing utilizes modality-specific mechanisms, then modality-specific phenomena should occur during conceptual processing. If variables such as occlusion, size, shape, and orientation affect perceptual processing, they should be likely to affect conceptual processing. Although not all perceptual variables should affect conceptual processing, at least some of them should. In contrast, if participants only use amodal representations during conceptual processing, then it is much less obvious that perceptual variables should affect performance. No amodal theory has ever predicted that variables like occlusion, size, shape, and orientation should affect conceptual processing.

3.1.3. Occlusion during property generation

Wu and Barsalou (2005) offer one example of this experimental logic. In several experiments, they assessed whether occlusion affects conceptual processing. If conceptual processing utilizes perceptual simulations, then a variable like occlusion, which affects vision, might affect conceptual processing as well. Wu and Barsalou manipulated occlusion by asking half the participants to generate properties for noun concepts (e.g., [LAWN]), and by asking the other half to generate properties for the same nouns preceded by revealing modifiers (e.g., [ROLLED-UP LAWN]). Wu and Barsalou predicted that if people simulate *LAWN* to generate its properties, they should rarely produce its occluded properties, such as DIRT and ROOTS. As in actual perception, occluded properties should not receive much attention, because they are hidden behind an object's surface. They also predicted that, conversely, when people produce properties for [ROLLED-UP LAWN], previously occluded properties would become salient in simulations and be produced more often.

Amodal theories of conceptual combination do not readily make this prediction, given that they do not anticipate effects of perceptual variables, such as occlusion. To the contrary, these theories typically assume that conceptual representations abstract over perceptual variation associated with occlusion, size, shape, and orientation. Furthermore, amodal theories of conceptual combination are relatively compositional in nature [e.g., Smith et al. (1988)]. When people combine ROLLED UP with [LAWN], for example, the meaning of the conceptual combination should be roughly the union of the individual meanings. Unless additional *post hoc* assumptions are added that produce interactions between nouns and modifiers, the properties for [LAWN] are not obviously changed by [ROLLED UP] (e.g., the accessibility of DIRT and ROOTS does not vary).

As Wu and Barsalou predicted, internal properties were produced relatively infrequently for the isolated nouns. Furthermore, the number of internal properties increased significantly when revealing modifiers were present, compared to when they were not present. Internal properties were also produced earlier in the protocols and in larger clusters. This finding occurred both for familiar noun combinations, such as *half watermelon*, and for novel ones, such as *glass car*. Rules for properties stored with the modifiers were not responsible for the increase in occluded properties (e.g., perhaps *rolled up* always increases the salience of unoccluded properties in the head noun). If such rules had been responsible, then a given modifier should have always increased the salience of normally occluded properties. However, in many noun phrases (e.g., *rolled-up snake*), this was not the case. Occluded properties only increased when the modifiers referred to entities whose internal parts become unoccluded in the process of conceptual combination (e.g., [ROLLED-UP LAWN]). Together, this pattern of results is consistent with the prediction that people construct simulations to represent conceptual combinations. When a simulation reveals occluded properties, they are produced more often, relative to when they remain occluded.

3.1.4. Size during property verification

Solomon and Barsalou (2004) assessed whether perceptual variables such as size affect the property verification task. Participants performed 200 property verification trials, half true and half false. Of primary interest was explaining the variance in response times (RTs) and error rates across the 100 true trials. Why are some concept–property pairs verified faster and with more accuracy than others? To assess this issue, the true concept–property pairs were scaled for perceptual, linguistic, and expectancy variables that might explain this variance. The linguistic variables included the associative strength between the concept and property words in both directions, the word frequency of the properties, and the word length. The perceptual variables included the size and position of the properties, whether they were occluded, whether they would be handled during situated action, and so forth. The expectancy variables assessed the polysemy of the property words (i.e., property words often have many different senses across objects; consider *leg, handle*).

The RTs and error rates for the 100 true trials were then analyzed by performing hierarchical linear regression on the three groups of variables. Most notably, the perceptual variables explained significant amounts of unique variance, after variance attributable to the linguistic and expectancy variables had been removed. Within the perceptual variables, the variable of size explained the most unique variance. As properties became larger, they took longer to verify.

This finding suggests that people verify properties by processing the regions of simulations that contain them. As the region that must be processed becomes larger, more time is required to process it. Kan et al. (2003) provided neural corroboration for this conclusion. When participants performed the Solomon and Barsalou experiment in a functional magnetic resonance imaging (fMRI) scanner, activation occurred in the left fusiform gyrus, an area often active in mental imagery and high-level object perception.

3.1.5. Shape during property verification

Solomon and Barsalou (2001) similarly found that detailed property shape was a critical factor in property verification. They assessed whether verifying a property facilitated verifying the same property again later for a different concept. For example, does verifying MANE for [LION] later facilitate verifying MANE for [PONY]? If participants represent MANE with a single amodal symbol that abstracts over differently shaped manes, then verifying MANE for [LION] should activate this symbol so that later, verifying MANE for [PONY] benefits from this. Alternatively, if people simulate manes to verify them, then simulating a lion mane might not later facilitate simulating a pony mane, given their differences in shape. Whereas a lion's mane wraps around the circumference of its neck, a pony's mane runs down the length of its neck. Because the shapes of the two manes differ significantly, simulating one might not facilitate simulating the other. Conversely, when the first mane verified has the same detailed shape as the later mane, facilitation should result (e.g., verifying MANE for [HORSE] prior to verifying MANE for [PONY]).

Across several experiments, the results supported the simulation view. When participants verified a property on an earlier trial, it facilitated verifying the same property later, but only if the detailed shape was similar. Thus, verifying MANE for [PONY] was facilitated by previously verifying MANE for [HORSE], but not by previously verifying MANE for [LION].

This effect did not result from greater overall similarity between [HORSE] and [PONY] than between [HORSE] and [LION]. When the property was highly similar for all three concepts, facilitation occurred from both the high and low similarity concepts. For example, verifying BELLY for [PONY] was facilitated as much by verifying BELLY for [LION] as by verifying BELLY for [HORSE]. Thus, the detailed perceptual similarity of the property was the critical factor, not the similarity between concepts. When the detailed shape of the critical property matched across two trials, facilitation occurred. When the detailed shapes differed, they did not. This pattern is again consistent with the conclusion that people simulate properties to verify them.

3.1.6. Modality switching during property verification

Further evidence for this conclusion comes from Pecher, Zeelenberg and Barsalou (2003), who found that a modality-switching phenomenon in perception also occurs during property verification. In actual perception, processing a signal on a modality suffers when the previous signal was perceived on a different modality than when it was perceived on the same modality [e.g., Spence, Nicholls and Driver (2000)]. For example, processing a light flash is faster when the previous signal was a light flash than when it was an auditory tone. A common explanation is that selective attention must shift to a new modality when the modality changes, incurring a temporal cost.

Pecher et al. demonstrated that the same phenomenon occurs during property verification (again using words for concepts and properties). When participants verified a conceptual property on one modality, processing was faster when the previous property came from the same modality than when it came from a different one. For example, verifying LOUD for [BLENDER] was faster when RUSTLING was verified for [LEAVES] on the previous trial than when TART was verified for [CRANBERRIES]. Analogous to perceptual modality switching, this switching cost suggests that people shift between modalities to simulate the properties being verified.

An alternative account is that properties from the same modality have higher associations between them than do properties from different modalities, which produce priming across adjacent trials. When associative strength was assessed, however, properties from the same modality were no more associated than were properties from different modalities. Furthermore, when highly associated properties were verified on contiguous trials in a later experiment, they were verified no faster than unassociated properties. Thus, modality switching and not associative strength appears to be responsible for the obtained effects.

Pecher, Zeehenberg and Barsalou (2004) further demonstrated the modality-switching effect under other task conditions. Marques (in press) extended the conditions that produce this effect, and also obtained this effect in experiments conducted in Portuguese. Barsalou et al. (2005) review the literature on modality switching in property verification. Again, it appears that the process of verifying properties produces simulations of them in modality-specific systems.

3.1.7. Shape and orientation during comprehension

All the evidence reviewed so far for simulation has come from conceptual tasks (i.e., property generation and property verification). Researchers, however, have also found evidence for simulation in language comprehension tasks. Because conceptual representations are widely assumed to underlie the representation of text meaning, these findings further implicate simulation in conceptual processing. The next several sections illustrate some of these recent findings.

Earlier we saw that property shape affects the process of verifying properties [Solomon and Barsalou (2001)]. Zwaan, Stanfield and Yaxley (2002) similarly found that shape affects language comprehension. Participants read a short vignette about an object that implied one of several possible shapes. Some participants read about a flying bird, whereas others read about a sitting bird (i.e., the implied shape of the bird's wings differed between the two vignettes). After reading the vignette, participants named a picture of an object in isolation, which was sometimes the same as an object just described in the previous sentence. On these trials, the shape of the object was manipulated such that it was either consistent or inconsistent with the implied shape in the sentence. When a bird was shown, for example, sometimes its wings were outstretched, and other times its wings were folded. As the simulation view predicts, participants named objects faster when the pictured shapes matched the implied shapes in the vignettes than when they did not.

Similarly, the implied orientation of an object affects comprehension. In Stanfield and Zwaan (2001), participants read vignettes that implied objects in particular orientations. For example, some participants read about someone pounding a nail into the wall, whereas other participants read about someone pounding a nail into the floor. Immediately afterwards, participants viewed a picture of an isolated object and had to indicate whether it had been mentioned in the vignette. Sometimes the orientation of an object that had occurred in the text matched its implied orientation, and sometimes it did not (e.g., a horizontal nail vs. a vertical one). Verification was fastest when the orientations matched. Again, participants appeared to be simulating objects mentioned in the sentences.

3.1.8. Movement direction in comprehension

In another set of comprehension studies, Glenberg and Kaschak (2002) found that people understood sentences that described actions by simulating the actions in their motor systems. In these experiments, participants read sentences and judged the grammaticality of the sentences. Embedded within the list of sentences were some that described actions moving toward the body (e.g., "Open the drawer.") vs. others that described actions moving away from the body (e.g., "Close the drawer."). Participants indicated that a sentence was grammatical by pressing a response button by moving either toward their bodies or away from their bodies. When the button press movements were consistent with the meaning of the sentence, RTs were faster than when they were inconsistent.

For example, participants were faster to verify that "Open the drawer." is grammatical with a button press toward their bodies than with a button press away from their bodies.

Glenberg and Kaschak also observed such effects for sentences that implied an abstract direction of motion. For example, participants were fastest to verify that "Liz told you the story" is grammatical with a button press toward their bodies. This pattern supports the conclusion that participants simulated the meanings of the sentences in their motor systems. When these simulations were consistent with response actions, processing was faster than when they were inconsistent.

3.1.9. Further evidence for simulation from comprehension studies

A variety of other findings further implicate simulation during text comprehension. Glenberg and Robertson (2000) found that readers readily compute the functional affordances of novel objects during comprehension, suggesting that they used perceptual simulations to represent them [also see Kaschak and Glenberg (2000)]. Fincher-Kiefer (2001) asked participants to adopt either a visual or verbal working memory load during text comprehension and found that the visual load produced the greatest interference on a subsequent inference task, suggesting that a simulated situation model represented the text. Richardson et al. (2003) found that the direction of a perceptual stimulus affected the time to process a sentence describing directional motion, suggesting that the sentences' meanings were being simulated. As these results indicate, accumulating evidence in the comprehension literature implicates simulation in the representation of text meaning.

3.1.10. Behavioral evidence for embodiment in social cognition

Perhaps the largest amount of evidence for simulation comes from social psychology. Because these findings have been reviewed in detail elsewhere, they are simply noted here [see Barsalou et al. (2003), Niedenthal et al. (in press)]. Many studies demonstrate that social stimuli induce bodily states. For example, perceiving various types of people and social events induces postures, arm movements, and facial expressions. The processing of a social stimulus does not simply activate an amodal description of the stimulus in memory. Instead, representations in the modalities become active to play central roles in social meaning.

Many other studies show that bodily states induce high-level, cognitive and emotional social representations. Specifically, various states of the body, arms, head, and face activate social categories, attitudes, and affects. Still other findings demonstrate that social information processing proceeds optimally when cognitive states are compatible with bodily states.

In general, the social psychology literature provides extensive evidence that social cognition is tightly coupled with the modalities, especially the motor, somatosensory, and limbic systems. States of these systems are often implicated in social processing, and have a variety of major effects upon it.

3.2. Neural evidence for a modal nonmodular conceptual system

As the previous section illustrated, a strong behavioral case has been developed for simulations in the conceptual system. A strong case has also been developed in the brain lesion and neuroimaging literatures.

3.2.1. Category-specific deficits

Lesions in a modality-specific system increase the likelihood of losing categories that rely on that system for processing exemplars. Because visual processing is important for interacting with [LIVING THINGS], such as [MAMMALS], damage to visual areas increases the chances of losing knowledge about these categories; because action is important for interacting with [MANIPULABLE OBJECTS], such as [TOOLS], damage to motor areas increases the chances of losing knowledge about these categories [e.g., Warrington and Shallice (1984), Warrington and McCarthy (1987), Damasio and Damasio (1994), Gainotti et al. (1995), Humphreys and Forde (2001)]. Similarly, lesions in color processing areas produce deficits in color knowledge [e.g., DeRenzi and Spinnler (1967)], and lesions in the spatial system produce deficits in location knowledge [e.g., Levine, Warach and Farah (1985)].

This pattern of findings has led many researchers to conclude that knowledge is grounded in the brain's modality-specific systems. Because the systems used to interact with a category's members during perception and action produce knowledge deficits when lesioned, category knowledge appears to rely on these systems for representational purposes.

Other factors besides damage to modality-specific systems also contribute to conceptual deficits. Caramazza and Shelton (1998) propose that localized brain areas represent specific categories that are evolutionarily important (e.g., [ANIMALS]). Tyler et al. (2000) propose that the statistical distribution of shared vs. unique property information for categories determines their vulnerability to lesion-based deficits. Thus, theories in this area increasingly include multiple mechanisms for explaining the variety of deficits observed [e.g., Coltheart et al. (1998), Cree and McRae (2003), Simmons and Barsalou (2003)]. Nevertheless, many researchers in this area have concluded that modality-specific systems play central roles in knowledge representation.

3.2.2. Neuroimaging studies of category knowledge

The neuroimaging literature further supports this conclusion [for reviews, see Pulvermüller (1999), Martin (2001)]. Consistent with the lesion literature, different types of categories differentially activate modality-specific systems. Categories that depend heavily on visual information (e.g., [ANIMALS]) strongly activate visual areas during neuroimaging, whereas categories that depend heavily on action (e.g., [TOOLS]) activate the motor system [e.g., Martin et al. (1996)]. Color categories activate color areas [e.g., Chao and Martin (1999)]. Social categories activate areas central to social interaction [e.g., Decety and Sommerville (2003), Gallese (2003)].

Consider several examples of these studies. Chao and Martin (2000) had participants view briefly presented pictures of manipulable objects, buildings, animals, and faces while lying passively in an fMRI scanner. While participants viewed manipulable objects (e.g., hammers), a brain circuit that underlies the grasping of manipulable objects became active. Notably, this circuit was not active while participants viewed buildings, animals, and faces. In previous studies with both monkey and humans, this grasping circuit became active either when participants actually performed actions with manipulable objects, or when they watched others perform such actions [e.g., Rizzolatti et al. (2002)]. Significantly, this circuit became active even though Chao and Martin's participants did not move in the scanner, and even though they did not view any agents or actions. Participants simply viewed pictures of static objects in isolation. Because the grasping circuit nevertheless became active, Chao and Martin concluded that this activation constituted a motor inference about how to act on the perceived object. Viewing a manipulable object activated category knowledge about it that included motor inferences (e.g., a hammer can be swung). Most importantly, simulations in the grasping circuit appeared to represent these inferences.

Simmons, Martin and Barsalon (in press) performed an analogous experiment with food categories. While participants lay passively in an fMRI scanner, they viewed food pictures for 2 s, and simply decided whether the current picture was identical to the previous one. Participants were not asked to categorize the foods, nor to think about how they taste. Nevertheless, under these superficial processing conditions, the pictures activated a brain area that represents how foods taste, along with areas that represent the reward value of foods. Even though participants were not actually tasting any foods, these areas became active. As participants perceived a food, it activated category knowledge, which then produced taste inferences via simulations in the gustatory system.

Modality-specific inferences from category knowledge also occur in response to words. Hauk, Johnsrude and Pulvermuller (2004) had participants simply read words for 2.5 s in an fMRI scanner. Within the list, randomly distributed subsets of words referred to head, arm, and leg actions (e.g., *lick*, *pick*, and *kick*, respectively). Hauk et al. predicted that if the meanings of action words are represented as simulations in the motor system, then all three types of words should activate it. Indeed, these words produced activations in motor areas that became active when participants actually moved the respective body parts in the scanner. Most notably, however, the three types of action words differentially activated their respective regions of the motor strip. Words for head actions, arm actions, and leg actions activated the regions that produce head, arm, and leg actions, respectively.

In summary, lesion and neuroimaging results from cognitive neuroscience corroborate the behavioral results reviewed earlier from cognitive and social psychology. These converging bodies of evidence support the first theme of this chapter: The conceptual system utilizes modality-specific mechanisms. Rather than being a modular system that only uses amodal representations, it is a nonmodular system that depends significantly on modal representations in modality-specific systems.

3.3. Evidence for situated conceptualizations

The second theme of this chapter is that conceptual representations are situated. Rather than being a general description of a category, a concept is the productive ability to generate many different situated conceptualizations, each supporting a different course of situated action with the category. As described earlier, a general description of a category used across all occasions would not provide the specialized inferences needed in particular situations. A single representation of [CHAIRS], for example, would be too general to produce the specific inferences needed to interact effectively with dining chairs, office chairs, theater chairs, airplane chairs, or ski lift chairs. Instead, each type of instance is best served by a situated conceptualization tailored to its respective situation.

Barsalou (2003b) proposed that a situated conceptualization supports situated action with a particular category member through four types of situated inferences:

(1) inferences about goal-relevant properties of the focal category;
(2) inferences about the background setting;
(3) inferences about likely actions that the agent could take to achieve an associated goal;
(4) inferences about likely introspective states that the agent might have while interacting with the category, such as evaluations, emotions, goals, and cognitive operations.

To see the importance of these four inferences types, imagine interacting with an airplane chair. Simply activating a general description of [CHAIRS] would provide insufficient inferences about these four aspects of situated action. A general [CHAIR] concept would not predict: (1) the particular parts of an airplane chair, (2) relevant aspects of the setting, (3) actions that could be performed on an airplane chair, and (4) introspections that might result.

Alternatively, activating a situated conceptualization supports all four types of prediction, including: (1) an airplane chair has controls for adjusting the seat back angle and headphone volume; (2) an airplane chair resides in a crowded setting with little space between adjacent chairs, so that tilting back one's chair impinges on the space of the passenger behind; (3) the action of pressing a light button on the chair activates an overhead reading light; (4) sitting in an airplane chair produces introspections that include negative affect about being in a cramped setting. As described earlier, these inferences are delivered via a pattern completion process that operates on situated conceptualizations. When a familiar category member is categorized, a situated conceptualization containing it becomes active and produces the four types of situated inferences.

The next four subsections review evidence that the conceptual system delivers these four types of situated inferences as people activate concepts. Rather than generic concepts becoming active, sets of situated inferences become active instead. Each of the next four subsections provides evidence in turn for one type of situated inference: goal-relevant properties of the focal category, background settings, actions, and introspective states.

3.3.1. Inferences about goal-relevant properties of the focal category

Many studies demonstrate that concepts do not produce the same generic representation over and over again across situations. Instead, a concept produces one of many possible representations tailored to the current context.

Barsalou (1982) illustrates this general finding. After participants had read a sentence, they verified whether a subsequent property was true or false of the subject noun. As the following examples illustrate, the predicate of the sentence varied between participants to manipulate the context of situated action:

The basketball was used when the boat sank.
The basketball was well worn from much use.

As can be seen, the predicate in each sentence situates [BASKETBALL] in a different context. Immediately after reading one of these two sentences, participants verified whether FLOATS was a true property of [BASKETBALL]. As the situated conceptualization view predicts, participants verified FLOATS 145 ms faster after reading the first sentence than after reading the second. As participants read about the boat sinking, the concept for [BASKETBALL] produced relevant inferences, such as FLOATS. The concept for [BASKETBALL] did not produce the same representation in both contexts.

Many additional studies across multiple literatures have found similar results [for a review, see Yeh and Barsalou (in press)]. In memory, context effects on word encoding are widespread [e.g., Greenspan (1986)]. In category learning, background knowledge about a situation constrains the properties of objects salient for them [e.g., Murphy (2000)]. In sentence processing, context effects on lexical access are legion [e.g., Kellas et al. (1991)]. Such findings clearly demonstrate that a general description of a category does not represent the category across situations. Instead, its representation is tailored to current task conditions.

3.3.2. Evidence for setting inferences

When the conceptual system represents a category, it does not do so in a vacuum. Instead, it situates the category in a background setting. As Yeh and Barsalou (in press) review, much work supports the inclusion of setting inferences in category representations.

Vallée-Tourangeau, Anthony and Austin (1998) provided one example of this evidence. On each trial, participants received the name of a category (e.g., [FRUIT]) and produced instances of it (e.g., [APPLE], [KIWI], [PEAR]). Accounts of this task typically assume that participants generate instances from conceptual taxonomies, or from similarity-based clusters. For [FRUIT], participants might first produce instances from [CITRUS FRUIT], then from [TROPICAL FRUIT], and then from [WINTER FRUIT]. Alternatively, Vallée-Tourangeau et al. proposed that participants situate the category in a background setting, scan across the setting, and report the instances present. To produce [FRUIT], for example, participants might imagine being in the produce section of their grocery store and report the instances found while scanning through it. Notably, this prediction assumes that categories are not represented in isolation. Instead, categories are associated with situations, such that categories and situations become active together in situated conceptualizations.

To assess this hypothesis, Vallée-Tourangeau et al. first asked participants to produce the instances of common taxonomic categories (e.g., [FRUIT]) and also of *ad hoc* categories (e.g., [THINGS PEOPLE TAKE TO A WEDDING]). After producing instances for all of these categories, participants were asked about the production strategies they had used, indicating one of three possible strategies for each category. First, if a category's instances came to mind automatically, participants indicated the *unmediated strategy*. Second, if the instances were accessed according to clusters in a taxonomy, participants indicated the *semantic strategy*. Third, if the instances were retrieved from experienced situations, participants indicated the *experiential strategy*.

As Vallée-Tourangeau et al. predicted, retrieving instances from situations was the dominant mode of production. Participants indicated they used the experiential strategy 54% of the time, followed by the semantic strategy (29%) and the unmediated strategy (17%). Surprisingly, this pattern occurred for both taxonomic and *ad hoc* categories. Because *ad hoc* categories arise in goal-directed situations, it is not surprising that they would be situated. Surprisingly, though, taxonomic categories were situated as well. Walker and Kintsch (1985) and Bucks (1998) reported similar findings.

The Wu and Barsalou (2005) experiments on occlusion described earlier further demonstrate that categories are situated in background settings. In those experiments, participants were explicitly instructed to produce properties of the target objects (e.g., [WATERMELON]). The instructions neither requested nor implied the relevance of background settings. Nevertheless, participants produced setting information regularly. Across experiments, the percentage of setting information ranged from 19% to 35%, averaging 25%. As participants simulated the target objects so that they could generate properties, they implicitly situated the objects in background settings, leading to the inadvertent production of many setting properties (e.g., PARK, PICNIC TABLE, etc., for [WATERMELON]).

Many further findings demonstrate a tight coupling between object representations and settings [as Yeh and Barsalou (in press) review in greater detail]. For example, studies of visual object processing have frequently shown that objects are strongly associated with their background scenes [e.g., Biederman (1981)]. On perceiving a familiar isolated object, a typical background scene is inferred immediately.

3.3.3. Evidence for action inferences

The previous two subsections illustrated that conceptual representations contain contextually relevant properties, and that they are situated in background settings. As next subsection illustrates, these situated representations of categories are not represented as detached from the conceptualizer. Instead, the conceptual system places the conceptualizer in these situated representations, producing inferences about possible actions the conceptualizer could take. In other words, the conceptual system implements simulations of "being there" with category members [Barsalou (2002)].

We have already seen a number of findings that support the presence of such inferences. As Glenberg and Kaschak (2002) demonstrated, reading a sentence about an action (with no agent mentioned) activates a motor representation of it. Rather than representing the action in a detached amodal manner, the action is represented as if the conceptualizer were preparing for situated action. The embodiment effects from social psychology mentioned earlier are also consistent with this conclusion. Relevant neuroimaging evidence was also reviewed. In Chao and Martin (2000), the grasping circuit became active when participants viewed manipulable artifacts in isolation. In Hauk et al. (2004), simply reading an action word activated the relevant part of the motor strip. All of these findings are consistent with the proposal that the conceptual system is action oriented – it is not simply a repository of amodal descriptions. Thus, when a concept becomes active, it prepares the conceptualizer for interacting with its instances by priming relevant actions in the motor system.

Adolphs et al. (2000) provide further evidence for this claim. On each trial, participants viewed a picture of face and indicated whether the expression was happy, sad, angry, etc. To assess the brain areas responsible for these categorizations, Adolphs et al. sampled from a registry of patients having lesions in different brain areas. To the extent that an area is important for categorizing visual expressions of emotion, lesions in the area should produce task deficits. The important finding was that large deficits resulted from lesions in the somatosensory cortex. Why would somatosensory lesions produce deficits on a visual task? Adolphs et al. proposed that simulating facial expressions on one's own face is central to visually recognizing expressions on other faces. When the somatosensory cortex is damaged, these simulations become difficult, and facial categorization suffers.

Work in social cognition corroborates this conclusion. Wallbott (1991) videotaped participants' own faces as they categorized emotional expressions on other's faces. Notably, participants simulated the emotional expressions that they were categorizing on their own faces. Furthermore, participants' accuracy was correlated with the extent to which their facial simulations were recognizable. Niedenthal et al. (2001) similarly found that preventing participants from simulating facial expressions decreased their ability to categorize other's facial expressions. Together, these findings illustrate that action systems in the brain become involved in the visual processing of faces.

Facial simulations of emotional expression in others can be viewed as motor inferences that the conceptual system produces to support situated action. Under many conditions, if another person is experiencing an emotion, it is often useful for the perceiver to adopt the same emotional state. Thus, if another person is happy about something, it is often supportive to be happy as well. Once the concept for a particular emotion becomes active, appropriate motor and somatosensory states for adopting it oneself follow as conceptual inferences.

3.3.4. Evidence for introspective state inferences

As we just saw, the conceptual system inserts the conceptualizer into situated conceptualizations via the simulation of possible actions in the situation. As this next subsection illustrates, the conceptual system further inserts the conceptualizer into situated conceptualizations via the simulation of possible introspections. As a particular situation associated with a category is simulated, relevant emotions, evaluations, goals, and cognitive operations are included.

Again, consider findings from the Wu and Barsalou (2005) occlusion experiments. Earlier we saw that participants produced properties about background settings, even though the explicit instruction was to produce properties of the target objects (e.g., [WATERMELON]). Of interest here is the fact that participants also produced many properties about the likely introspective states that they would experience in associated situations. These included evaluations of whether objects are good, bad, effective, ineffective, etc. They also included emotional reactions to objects, such as happiness, along with other cognitive operations relevant to interacting with them (e.g., comparing an object to alternatives). On the average, 10% of the properties that participants produced were about introspective states, ranging from 6% to 15% across experiments. As participants simulated a target object, they situated it in a background setting and included themselves as agents. In the process, they simulated likely introspective states that they would experience, which they then inadvertently described in their protocols.

Adopting particular perspectives in laboratory tasks further implicates introspective states in conceptual simulation. A perspective can be viewed as an introspective state because it reflects one of many possible views that an agent could take in a physical situation. As people simulate situations, they imagine likely perspectives. In Spivey et al. (2000), participants listened to a vignette while wearing an eye-tracking helmet (which they believed was turned off at the time). Whereas some participants heard about a skyscraper, others heard about a canyon. As participants listened to the vignette, they tended to adopt the relevant perceptual perspective on the setting. Participants hearing about the skyscraper were most likely to look up, whereas participants hearing about the canyon were most likely to look down. Participants acted as if they were "there," adopting the relevant perspective on the situation.

In a related study, Barsalou and Barbey (2005) videotaped participants as they produced properties for object concepts. On a few trials, the object was something that would typically be encountered above a person (e.g., [BIRD]) or on the ground (e.g., [WORM]). When participants produced properties for objects typically found above them, their eyes, face, and hands were more likely to drift up than for objects typically found below them, and vice versa. Again, participants appeared to simulate the perspective on the situation that they would take if interacting with the object.

3.3.5. Evidence for dynamical simulations

As we have seen, a general decontextualized description does not represent a category across situations. Instead, conceptual representations appear to be highly contextualized, containing various types of specialized inferences that support situated action in specific contexts. If this account is correct, then another prediction follows: Many different representations of a category should be observed both between different individuals and in one individual. When different individuals represent a given category, their representations of it should differ, depending on the situation that they are anticipating. Similarly, when the same individual represents a category on different occasions, its representations should again differ, depending on the anticipated situation.

A variety of findings support this prediction [as reviewed in Barsalou (1987, 1989, 1993)]. Consider the agreement for typicality judgments across different members of a category (e.g., the typicality of different birds). Across a variety of studies, the average correlation in typicality judgments between pairs of participants for a given category was only around 0.40. Different participants appeared to use very different prototypes for judging typicality. Additionally, individual participants appeared to use different category representations on different occasions. When the same participant judged typicality 2 weeks after an initial judgment, the average correlation with their earlier judgment was about 0.80, indicating a change in how they represented the categories.

Similar variability arises in other tasks. In McCloskey and Glucksberg (1978), participants assigned basic-level categories to superordinates in two sessions separated by a month. Across the two sessions, roughly 25% of the basic-level categories changed superordinate membership, indicating variability in categorization criteria. In Barsalou (1989), participants exhibited variability in property generation. On average, two participants only produced 44% of the same properties for the same category. Over a two-week delay, the same participant only produced 66% of the same properties.

On the basis of these results, Barsalou (1987, 1989, 1993) concluded that a concept is a dynamic system. Depending on a variety of conditions, a concept produces a wide variety of different representations. After reviewing similar findings from the conceptual development literature, Smith and Samuelson (1997) reached a similar conclusion. Together, these results challenge the view that a concept is a general description used over and over again across situations. Instead, a concept appears to be an ability or skill to construct specific representations that support different courses of situated action. Because a concept produces a wide variety of situated conceptualizations, substantial variability in its representation arises.

4. Conclusion

Increasing empirical evidence supports the two themes of this chapter. First, the conceptual system does not appear to be modular, nor to solely use amodal representations. Instead, it appears to share mechanisms with modality-specific systems, such that its representations are often modal. Second, the conceptual system does not traffic solely in general descriptions of categories that remain stable across contexts. Instead, it dynamically produces contextualized representations that support situated action in different situations.

The construct of a situated conceptualization integrates these two themes, where a situated conceptualization is a multimodal simulation that supports one specific course of goal-directed interaction with a particular category instance. A given concept produces many different situated conceptualizations, each tailored to a different instance and setting. A situated conceptualization creates the experience of "being there" with a category instance via integrated simulations of agents, objects, settings, actions, and introspections. On recognizing a familiar type of category instance, an entrenched situated conceptualization associated with it becomes active to provide relevant inferences via pattern completion.

4.1. Important issues for future research

Clearly, many outstanding issues remain, including the role of amodal symbols in the conceptual system, the implementation of classic symbolic functions, and the representation of abstract concepts. Each is addressed briefly in turn.

4.1.1. Amodal symbols

One obvious question is whether amodal symbols coexist with simulations in the conceptual system. Perhaps amodal symbols are necessary to perform classic symbolic functions that simulations cannot implement. Barsalou (1999b, 2003a), however, argues that a simulation-based system can, at least in principle, implement these functions. It remains an open empirical question how these functions are implemented, although some evidence suggests that simulations contribute to them (as discussed shortly).

One possibility is that the classic amodal symbols found in predicate calculus-based approaches to representation do not exist in the conceptual system, but that other types of (relatively) amodal symbols do. For example, the sets of conjunctive neurons in association areas that trigger simulations in feature systems could be viewed as amodal vectors [Simmons and Barsalou (2003)]. Damasio (1989), however, argues that conjunctive neurons serve only to control simulations – they do not function as stand-alone representations. Simmons and Barsalou (2003) further propose that conjunctive neurons have modality-specific tunings, such that they are not truly amodal.

Another obvious issue is whether all conceptual representations are situated. Although the focus here has been on situated conceptualizations, abstractions appear to occur ubiquitously throughout conceptual processing. During everyday cognition, people appear to frequently process generalizations about categories, such as "Hondas are reliable," "Reality TV shows are boring," and "Academics are liberal." Barsalou (2003a) argues that such abstractions are central to human cognition, and that it is essential for simulation theories to explain them. In this spirit, Barsalou (2003a) proposes that simulators implement abstractions via productively constructed configurations of simulations. Much future work is clearly needed to determine whether complex simulations do indeed underlie abstractions. It is also important to establish how abstractions and situated conceptualizations interact during conceptual processing.

4.1.2. Symbolic functions

Another important issue concerns whether simulations implement classic symbolic functions, such as the type-token distinction, categorical inference, productivity, and propositions. Barsalou (1999b, 2003a) argues that, at least in principle, simulations can implement these functions. Preliminary empirical evidence suggests that simulations are involved to some extent. For example, evidence reviewed earlier suggests that the process of predicating a property of a category – a central symbolic function – relies on simulations. When people perform the property verification task, perceptual variables such as size and shape affect performance [Solomon and Barsalou (2001, 2004)]. The time to verify a property becomes shorter the smaller a property is and the more its shape matches property simulations run earlier. If people verify properties by processing regions of simulations, size and shape would have these sorts of effects.

We also saw evidence earlier that simulations underlie conceptual combination, another central symbolic function. Specifically, Wu and Barsalou (2005) found that occlusion affected the properties produced for conceptual combinations. When properties were occluded in isolated nouns, they were produced infrequently. When properties were unoccluded in conceptual combinations, their production increased. This finding suggests that simulations are used to combine modifier and noun concepts.

Clearly, much more work is necessary to establish how the brain implements classic symbolic functions. Preliminary evidence, though, suggests that simulations play a role.

4.1.3. Abstract concepts

Many people find it intuitively easier to understand how the simulation view explains concrete concepts than how it explains abstract concepts. Barsalou (1999b), however, proposed that the conceptual content of abstract concepts is drawn from events and introspections, and that this content could be simulated in relevant modality-specific systems. Rather than being redescribed with amodal symbols, the content of abstract concepts could be simulated just as for concrete concepts, with the difference being in the content simulated.

Barsalou and Wiemer-Hastings (2005) provided preliminary evidence for this view. When participants generated properties for abstract and concrete concepts, they tended to produce situated content for both, including agents, objects, settings, actions, and introspections. Indeed, the distributions of content were remarkably similar. As predicted, the difference was that the concrete concepts contained more content about physical objects, background settings, and simple behaviors, whereas the abstract concepts contained more content about people, social interactions, complex relations, and introspections. In general, the content reported for abstract concepts was drawn from the four types of content for situated conceptualizations reviewed earlier. Thus, it appears possible that this content could be simulated.

Clearly, much more work on abstract concepts is needed to understand how the brain represents this fundamentally important type of concept. Nevertheless, simulation may play a central role.

Acknowledgment

This work was supported by National Science Foundation Grants SBR–9421326, SBR–9796200, SBR–9905024, and BCS–0212134 to Lawrence W. Barsalou.

References

Adolphs, R., H. Damasio, D. Tranel, G. Cooper, and A.R. Damasio (2000), "A role for somatosensory cortices in the visual recognition of emotion as revealed by three-dimensional lesion mapping", Journal of Neuroscience 20:2683–2690.

Andersen, S.M., and S. Chen (2002), "The relational self: an interpersonal social-cognitive theory", Psychological Review 109:619–645.

Barsalou, L.W. (1982). "Context-independent and context-dependent information in concepts", Memory & Cognition 10:82–93.

Barsalou, L.W. (1987), "The instability of graded structure: implications for the nature of concepts", in: U. Neisser, ed., Concepts and conceptual development: Ecological and Intellectual Factors in Categorization (Cambridge University Press, Cambridge), 101–140.

Barsalou, L.W. (1989), "Intraconcept similarity and its implications for interconcept similarity", in: S. Vosniadou and A. Ortony, eds., Similarity and Analogical Reasoning. (Cambridge University Press, Cambridge) 76–121.

Barsalou, L.W. (1993), "Structure, flexibility, and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols", in: A.C. Collins, S.E. Gathercole and M.A. Conway, eds., Theories of Memory (Lawrence Erlbaum Associates, London), 29–101.

Barsalou, L.W. (1999a), "Language comprehension: archival memory or preparation for situated action", Discourse Processes: 28:61–80.

Barsalou, L.W. (1999b), "Perceptual symbol systems", Behavioral and Brain Sciences 22: 577–660.

Barsalou, L.W. (2002), "Being there conceptually: simulating categories in preparation for situated action", in: N.L. Stein, P.J. Bauer and M. Rabinowitz , eds., Representation, Memory, and Development: Essays in Honor of Jean Mandler (Erlbaum, Mahwah, NJ) 1–16.

Barsalou, L.W. (2003a), "Abstraction in perceptual symbol systems", Philosophical Transactions of the Royal Society of London: Biological Sciences 358:1177–1187.

Barsalou, L.W. (2003b), "Situated simulation in the human conceptual system", Language and Cognitive Processes 18:513–562.

Barsalou, L.W., and A. Barbey (2005), "Spontaneous body movements during property generation for concepts", in preparation.

Barsalou, L.W., P.M. Niedenthal, A. Barbey and J. Ruppert (2003), "Social embodiment", in: B. Ross, ed., The Psychology of Learning and Motivation, vol. 43 (Academic Press, San Diego) 43–92.

Barsalou, L.W., D. Pecher, R. Zeelenberg, W.K. Simmons and S.B. Hamann (2005), Multi-modal simulation in conceptual processing", in: W. Ahn, R. Goldstone, B. Love, A. Markman, and P. Wolff eds., Categorization Inside and Outside the Lab: Essays in Honor of Douglas L. Medin (American Psychological Association, Washington, DC), 249–270.

Barsalou, L.W., and K. Wiemer-Hastings (2005), "Situating abstract concepts", in: D. Pecher and R. Zwaan, eds., Grounding Cognition: The role of Perception and Action in Memory, Language, and Thought. (Cambridge University Press, New York), 129–163.

Biederman, I. (1981), "On the semantics of a glance at a scene", in: M. Kubovy and J.R. Pomerantz, eds., Perceptual Organization. (Lawrence Erlbaum Associates, Hillsdale, NJ). 213–253.

Brewer, W.F., and J. C. Treyens (1981), "Role of schemata in memory for places", Cognitive Psychology 13:207–230.

Bruner, J.S. (1957), "Going beyond the information given", in: J.S. Bruner, E, Brunswik, L. Festinger, F. Heider, K.F. Muenzinger, C.E. Osgood and D. Rapaport, eds., Contemporary Approaches to Cognition. (Harvard University Press, Cambridge, MA). 41–69.

Bucks, R.S. (1998), "Intrusion errors in Alzheimer's disease", Doctoral Dissertation (University of Bristol).

Caramazza, A., and J.R. Shelton (1998), "Domain-specific knowledge systems in the brain: the animate-inanimate distinction", Journal of Cognitive Neuroscience 10:1–34.

Chao, L.L., and A. Martin (1999), "Cortical regions associated with perceiving, naming, and knowing about colors", Journal of Cognitive Neuroscience 11:25–35.

Chao L.L., and A. Martin (2000), "Representation of manipulable man-made objects in the dorsal stream", Neuroimage 12:478–84.

Collins, A.M., and E.F. Loftus (1975), "A spreading activation theory of semantic processing", Psychological Review 82:407–428.

Collins, A.M., and M.R. Quillian (1969), "Retrieval time from semantic memory", Journal of Verbal Learning and Verbal Behavior 8:240–248.

Coltheart, M., L. Inglis, L. Cupples, P. Michie, A. Bates and B. Budd (1998), "A semantic subsystem of visual attributes", Neurocase 4(4):353–370.

Cree, G.S., and K. McRae (2003), "Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns)", Journal of Experimental Psychology: General 132:163–201.

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

Damasio, A.R., and H. Damasio (1994), "Cortical systems for retrieval of concrete knowledge: the convergence zone framework", in: C. Koch and J.L. Davis, eds., Large-scale neuronal theories of the brain (MIT Press, Cambridge, MA) 61–74.

Decety, J., and J.A. Sommerville (2003), "Shared representations between self and other: a social cognitive neuroscience view", Trends in Cognitive Sciences 12:527–533.

DeRenzi, E., and H. Spinnler (1967), "Impaired performance on color tasks in patients with hemispheric lesions", Cortex 3:194–217.

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

Farah, M.J., and J.L. McClelland (1991), "A computational model of semantic memory impairment: modality specificity and emergent category specificity", Journal of Experimental Psychology: General 120:339–357.

Fincher-Kiefer, R. (2001), "Perceptual components of situation models", Memory and Cognition 29:336–343.

Finke, R.A. (1989), Principles of Mental Imagery. (MIT Press,Cambridge, MA).

Gainotti, G., M.C. Silveri, A. Daniele and L. Giustolisi (1995), "Neuroanatomical correlates of category-specific semantic disorders: a critical survey", Memory 3:247–264.

Gallese V. (2003), "The roots of empathy: the shared manifold hypothesis and the neural basis of intersubjectivity", Psychopathology 36:171–180.

Glenberg, A.M., and M.P. Kaschak (2002), "Grounding language in action", Psychonomic Bulletin & Review 9:558–569.

Glenberg, A.M., and D.A. Robertson (2000), "Symbol grounding and meaning: a comparison of high-dimensional and embodied theories of meaning", Journal of Memory and Language 43:379–401.

Greenspan, S.L. (1986), "Semantic flexibility and referential specificity of concrete nouns", Journal of Memory and Language 25:539–557.

Grezes, J., and J. Decety (2001), "Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta-analysis", Human Brain Mapping 12:1–19.

Hampton, J.A. (1979), "Polymorphous concepts in semantic memory", Journal of Verbal Learning and Verbal Behavior 18:441–461.

Hauk, O., I. Johnsrude and F. Pulvermüller (2004). Somatopic representation of action words in human motor and premotor cortex. Neuron, 41:301–307.

Humphreys, G.W., and E.M.E. Forde (2001), "Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits", Behavioral & Brain Sciences 24:453–509.

Johnson-Laird, P.N. (1983), Mental models. Cambridge, MA: Harvard University Press.

Kan, I.P., L.W., Barsalou, K.O. Solomon, J.K. Minor and S.L. Thompson-Schill (2003), "Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge", Cognitive Neuropsychology 20:525–540.

Kaschak, M.P., and A.M. Glenberg (2000), "Constructing meaning: the role of affordances and grammatical constructions in sentence comprehension", Journal of Memory and Language 43:508–529.

Kellas, G., S.T. Paul, M. Martin and G.B. Simpson (1991), "Contextual feature activation and meaning access", in: G.B. Simpson ed., Understanding word and sentence (Elsevier, New York).

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.

Kosslyn, S.M. (1994), Image and brain. (MIT Press,Cambridge, MA).

Levine, D.N., J. Warach and M.J. Farah (1985), "Two visual systems in menal imagery: dissociation of "What" and "Where" in imagery disorders due to bilaterial posterior cerebral lesions", Neurology 35:1010–1018.

Marques, J.F. (in press), "Specialization and semantic organization: evidence for multiple-semantics linked to sensory modalities", Memory & Cognition.

Martin, A. (2001), "Functional neuroimaging of semantic memory", in: R. Cabeza and A. Kingstone, eds., Handbook of Functional Neuroimaging of cognition (MIT Press,Cambridge, MA) 153–186.

Martin, A., C.L. Wiggs, L.G. Ungerleider and J.V. Haxby (1996), "Neural correlates of category-specific knowledge", Nature 379:649–652.

McCloskey, M., and Glucksberg, S. (1978), "Natural categories: well-defined or fuzzy sets", Memory & Cognition 6:462–472.

McRae, K., V.R. de Sa and M.S. Seidenberg (1997), "On the nature and scope of featural representations of word meaning", Journal of Experimental Psychology: General 126:99–130.

Murphy, G.L. (2000), "Explanatory concepts", in: R.A. Wilson and F.C. Keil, eds., Explanation and Cognition. (MIT Press, Cambridge, MA) 361–392.

Murphy, G.L., and D.L. Medin, (1985), "The role of theories in conceptual coherence", Psychological Review 92:289–316.

Niedenthal, P.M., L.W. Barsalou, P. Winkielman, S. Krauth-Gruber and F. Ric (in press), "Embodiment in attitudes, social perception, and emotion", Personality and Social Psychology Review.

Niedenthal, P.M., M. Brauer, J.B. Halberstadt and A.H. Innes-Ker (2001), When did her smile drop? Facial mimicry and the influences of emotional state on the detection of change in emotional expression. Cognition and Emotion 15:853–864.

Palmer, S.E. (1999), Vision Science: Photons to Phenomenology (MIT Press, Cambridge, MA).

Pecher, D., R. Zeelenberg and L.W. Barsalou (2003), "Verifying properties from different modalities for concepts produces switching costs", Psychological Science 14:119–124.

Pecher, D., R. Zeelenberg, and L.W. Barsalou (2004), "Sensorimotor simulations underlie conceptual representations: modality-specific effects of prior activation", Psychonomic Bulletin & Review 11:164–167.

Pulvermüller, F. (1999), "Words in the brain's language", Behavioral and Brain Sciences 22:253–336.

Richardson, D.C., M.J. Spivey, L.W. Barsalou and K. McRae (2003), "Spatial representations activated during real-time comprehension of verbs", Cognitive Science 27: 767–780.

Rizzolatti, G., L. Fadiga, L. Fogassi and V. Gallese (2002), "From mirror neurons to imitation: facts and speculations", in: A.N. Meltzoff and W. Prinz eds., The Imitative mind: Development, Evolution, and Brain Bases. Cambridge Studies in Cognitive Perceptual Development. (Cambridge University Press, New York) 247–266.

Rosch, E., and C.B. Mervis (1975), "Family resemblances: studies in the internal structure of categories", Cognitive Psychology 7:573–605.

Schyns, P.G., R.L. Goldstone and J.P. Thibaut (1998), "The development of features in object concepts", Behavioral and Brain Sciences 21:1–54.

Simmons, K., and L.W. Barsalou (2003), "The similarity-in-topography principle: reconciling theories of conceptual deficits", Cognitive Neuropsychology 20:451–486.

Simmons, W.K., A. Martin and L.W. Barsalou (in press), "Pictures of appetizing foods activate gustatory cortices for taste and reward", Cerebral Cortex.

Smith, E.E. (1978), "Theories of semantic memory", in: W.K. Estes, ed., Handbook of Learning and Cognitive Processes vol. 6, (Lawrence Erlbaum Associates, Hillsdale, NJ) 1–56.

Smith, E.E., D.N. Osherson, L.J. Rips and M. Keane (1988), "Combining prototypes: a selective modification model", Cognitive Science 12: 485–528.

Smith, L.B., and L.K. Samuelson (1997), "Perceiving and remembering: category stability, variability and development", in: K. Lamberts and D.R. Shanks, ed., Knowledge, Concepts and Categories. Studies in Cognition. (MIT Press, Cambridge, MA). 161–195.

Solomon, K.O., and L.W. Barsalou (2001), "Representing properties locally", Cognitive Psychology 43:129–169.

Solomon, K.O., and L.W. Barsalou (2004), "Perceptual simulation in property verification", Memory & Cognition 32:244–259.

Spence, C., M.E.R. Nicholls and J. Driver (2000), "The cost of expecting events in the wrong sensory modality". Perception & Psychophysics 63:330–336.

Spivey, M., M. Tyler, D. Richardson and E. Young (2000), "Eye movements during comprehension of spoken scene descriptions", Proceedings of the 22nd Annual Conference of the Cognitive Science Society (Erlbaum, Mahwah, NJ) 487–492

Stanfield, R.A., and R.A. Zwaan (2001), "The effect of implied orientation derived from verbal context on picture recognition", Psychological Science 12:153–156.

Tulving, E. (1972), "Episodic and semantic memory", in: E. Tulving and W. Donaldson, eds., Organization and Memory. (Academic, New York) 381–403

Tyler, L.K., H.E. Moss, M.R. Durrant-Peatfield and J.P. Levy (2000), "Conceptual structure and the structure of concepts: a distributed account of category-specific deficits", Brain & Language 75:195–231.

Vallée-Tourangeau, F., S.H. Anthony and N.G. Austin (1998), "Strategies for generating multiple instances of common and ad hoc categories", Memory 6:555–592.

Walker, W.H., and W. Kintsch (1985), "Automatic and strategic aspects of knowledge retrieval", Cognitive Science 9:261–283.

Wallbott, H.G. (1991), "Recognition of emotion from facial expression via imitation? Some indirect evidence for an old theory", British Journal of Social Psychology 30:207–219.

Warrington, E.K., and R.A. McCarthy (1987), "Categories of knowledge: further fractionations and an attempted integration", Brain 110:1273–1296.

Warrington, E.K., and T. Shallice (1984), "Category specific semantic impairments", Brain 107: 829–854.

Wiemer-Hastings, K., J. Krug and X. Xu (2001), "Imagery, context availability, contextual constraint, and abstractness", Proceedings of the 23rd Annual Conference of the Cognitive Science Society (Erlbaum, Mahwah, NJ) 1134–1139.

Wu, L., and L.W. Barsalou (2005), "Perceptual simulation in property generation", under review.

Yeh, W., and L.W. Barsalou (in press), "The situated character of concepts", American Journal of Psychology.

Zatorre, R.J., A.R. Halpern, D.W. Perry, E. Meyer and A.C. Evans (1996), "Hearing in the mind's ear: a PET investigation of musical imagery and perception", Journal of Cognitive Neuroscience 8:29–46.

Zeki, S. (1993), A Vision of the Brain. (Blackwell Scientific, Oxford).

Zwaan, R.A., and G.A. Radvansky (1998), Situation models in language comprehension and memory. Psychological Bulletin 123:162–185.

Zwaan, R.A., R.A. Stanfield and R.H. Yaxley (2002), "Language comprehenders mentally represent the shapes of objects", Psychological Science 13:168–171.