

Cognitively Plausible Theories of Concept Composition

Lawrence W. Barsalou

Institute of Neuroscience and Psychology
School of Psychology
University of Glasgow

29 July 2015

Chapter to appear in:

Y. Winter & J. A. Hampton (Eds.), *Compositionality and concepts in linguistics and psychology*.
London: Springer Publishing.

The human ability to construct and combine concepts is prolific. On the one hand, people acquire tens of thousands of concepts for diverse categories of settings, agents, objects, actions, mental states, bodily states, properties, relations, and so forth. On the other, people combine these concepts to construct infinite numbers of more complex concepts, as the open-ended phrases, sentences, and texts that humans produce effortlessly and ubiquitously illustrate. Major changes in the brain, the emergence of language, and new capacities for social cognition all probably played central roles in the evolution of these impressive conceptual abilities (e.g., Deacon, 1997; Donald, 1993; Tomasello, 2009).

In psychology alone, much research addresses human concepts (e.g., Barsalou, 2012; Murphy, 2002; Smith & Medin, 1981) and concept composition (often referred to as *conceptual combination*; e.g., Costello & Keane, 2000; Gagné & Spalding, 2014; Hampton, 1997; Hampton & Jönsson, 2012; Medin & Shoben, 1988; Murphy, 1988; Wisniewski, 1997; Wu & Barsalou, 2009). More generally across the cognitive sciences, much additional research addresses concepts and the broader construct of compositionality (for a recent collection, see Werning, Hinzen, & Machery, 2012).

Background Framework

A grounded approach to concepts. Here I assume that a concept is a dynamical distributed network in the brain coupled with a category in the environment or experience, with this network guiding situated interactions with the category's instances (for further detail, see Barsalou, 2003b, 2009, 2012; in press-a, in press-b). The concept of

bicycle, for example, represents and guides interactions with the category of bicycles in the world. Across interactions with a category's instances, a concept develops in memory by aggregating information from perception, action, and internal states. Thus, the concept of *bicycle* develops from aggregating multimodal information related to bicycles across the situations in which they are experienced. As a consequence of using selective attention to extract information relevant to the concept of *bicycle* from the current situation (e.g., a perceived bicycle), and then using integration mechanisms to integrate it with other bicycle information already in memory, aggregate information for the category develops continually (Barsalou, 1999). As described later, however, background situational knowledge is also captured that plays important roles in conceptual processing (Barsalou, in press-b, 2003b; Yeh & Barsalou, 2006). Although learning plays central roles in establishing concepts, genetic and epigenetic processes constrain the features that can be represented for a concept, and also their integration in the brain's association areas (e.g., Simmons & Barsalou, 2003). For example, biologically-based neural circuits may anticipate the conceptual structure of evolutionarily important concepts, such as agents, minds, animals, foods, and tools.

Once the conceptual system is in place, it supports virtually all other forms of cognitive activity, both online in the current situation and offline when representing the world in language, memory, and thought (e.g., Barsalou, 2012, in press-a, in press-b). From the perspective developed here, when conceptual knowledge is needed for a task, concepts produce situation-specific simulations of the relevant category

dynamically, where a simulation attempts to reenact the kind of neural and bodily states associated with processing the category. On needing conceptual knowledge about *bicycles*, for example, a small subset of the distributed bicycle network in the brain becomes active to simulate what it would be like to interact with an actual bicycle. This multimodal simulation provides anticipatory inferences about what is likely to be perceived further for the bicycle in the current situation, how to interact with it effectively, and what sorts of internal states might result (Barsalou, 2009). The specific bicycle simulation that becomes active is one of infinitely many simulations that could be constructed dynamically from the bicycle network—the entire network never becomes fully active. Typically, simulations remain unconscious, at least to a large extent, while causally influencing cognition, affect, and action. To the extent that part of a simulation becomes conscious, mental imagery is experienced. Such simulations need not provide complete or accurate representations, but are likely to be incomplete and distorted, representing abstractions, caricatures, and ideals, as well as specific learning episodes.

Situating concepts. As Pelletier (this volume) notes, theories of concepts often focus either on concepts in the world from the formalist perspective (e.g., truth conditions) or on concepts in the brain from the cognitivist perspective (what Pelletier refers to as *objectivist* vs. *subjectivist* approaches). Although the grounded approach to concepts makes many assumptions about concepts from the cognitivist perspective, it additionally assumes that physical situations play central roles in establishing and using concepts (e.g., Barsalou, 2003b, in press-b; Barsalou et al., 1993; Barsalou, Breazeal, & Smith, 2007). From the grounded perspective, concepts are typically situated, that is, they become active to process some aspect of the current physical situation (or a physical situation in the past or future). Although concepts may sometimes become active independently of a situation, they typically become active to support effective action in a specific situation. As a result, concepts become coupled with their physical referents. Although grounded theories of concepts have a long ways to go in developing satisfactory accounts of this coupling, they naturally anticipate it, and provide many mechanisms for understanding and explaining it. Thus, grounded theories offer an approach for unifying formalist and cognitivist accounts of concepts, and further assume that neither approach alone is sufficient.

Overview. The grounded account of concepts

just described is not all that unusual given the diverse theories of conceptual processing that exist within the prototype, exemplar, connectionist, Bayesian, psycholinguistic, and grounded traditions of cognitive science (e.g., Asher, 2011; Barsalou, 1990; Bosch, 1983; Casasanto & Lupyan, in press; Connell & Lynott, 2014; Evans, 2009; Jones & Love, 2011; Lebois, Wilson-Mendenhall, & Barsalou, 2014; McClelland & Rumelhart, 1985; Minda & Smith, 2011; Pothos & Wills, 2011). Many cognitively plausible theories assume increasingly that concepts are represented dynamically and are adapted to specific situations.

In contrast, theories of concept composition often assume relatively idealized accounts of concepts, where much of the dynamically varying content of concepts is omitted, given its complexity and messiness, focusing instead on a few key features and relations (e.g., some of the chapters in Werning et al., 2012). Even in the cognitivist accounts of concept composition mentioned earlier, concepts are often relatively idealized. Idealizing concepts has made establishing principles of concept composition tractable, and doing so has clearly been an effective strategy.

Nevertheless, avoiding the complexities of conceptual processing may obscure important principles of how concepts combine. As theories of concept composition mature increasingly, theorists and researchers appear willing to take on the challenge of dealing with these complexities (as illustrated by many of the chapters in this volume and also in Werning et al., 2012). In this spirit, I focus on three principles of human conceptual processing that strike me as central to its character: content variability, multiple representational forms, and pragmatic constraint. On the one hand, these principles offer significant challenges to theories of concept composition. On the other, seriously addressing and incorporating these challenges may yield new theories that provide increasingly powerful and insightful accounts.

Content Variability

Lack of conceptual cores. Researchers often assume that a concept has a *core*, namely, important information about the respective category that is consistently activated rapidly and automatically, independently of context. Although cores are sometimes assumed to contain rules or definitions (e.g., “unmarried adult male human” for *bachelors*), they can also contain typical features that are neither necessary nor sufficient for category membership (e.g., “small flying feathered singing animal” for *birds*). Not only have

traditional accounts of concepts proposed that concepts have cores (e.g., Barsalou, 1982, 1989; Conrad, 1978; Greenspan, 1986; Whitney, McKay, Kellas, & Emerson, 1985), so have current accounts (e.g., Dove, 2009; Machery, 2007; Mahon & Caramazza, 2008).

An interesting wrinkle to core theories is whether cores become active early vs. late (Lebois et al., 2014). According to most core theories, including those just cited, a conceptual core is the first information active for a concept, given its importance in representing the respective category (i.e., “core early” theories). Interestingly, however, major accounts of concepts over the years have assumed the opposite, namely, conceptual cores are only activated late, if and when needed (i.e., “core late” theories). Rather than being automatic and context-dependent, cores tend to be slow and optional.

In some core late theories, definitional information is considered relatively specialized and only necessary when careful judgments about a category must be made (e.g., Keil & Batterman, 1984; Malt, 1994; Smith, Shoben, & Rips, 1974). According to these theories, typical non-core information is sufficient most of the time for processing a category. When processing *water*, for example, it is usually sufficient to represent it as a colorless, clear, liquid used for daily needs, and widely available from faucets. Only on rare occasions does it become necessary to access and use core definitional information that water is H₂O (Malt, 1994).

In other core late theories, core information tends to be active late because typical information is easier to activate, thus becoming active initially (e.g., Blanchette & Dunbar, 2000; Forbus, Gentner, & Law, 1995; Gentner, Rattermann, & Forbus, 1993; Ross, 1984, 1987). From this perspective, cues first retrieve typical information from concepts related to superficial properties of the situation. Even though deeper knowledge might be available and relevant, it tends to be displaced by relatively superficial information that becomes active more easily. Only when strategic efforts to activate deeper, more important information are made does this latter information become available. Importantly for our purposes, however, core information in core late theories is not activated quickly and automatically across contexts.

Even more problematically for core early views, many theories argue that concepts don't have cores at all (e.g., Casasanto & Lupyan, in press; Connell & Lynott, 2014; Evans, 2009; McCloskey & Glucksberg, 1979). According to

these theories, all information in a concept is context-dependent, at least to some extent. Although some information for a concept may be highly entrenched (e.g., feathers for *birds*), it is still modulated by context. Indeed much recent research suggests that central information in concepts isn't activated automatically across contexts (for reviews, see Gawronski & Cesario, 2013; Kiefer et al., 2012; Lebois et al., 2014; but see Augustinova & Ferrand, 2014). For example, in the Stroop task that has long been believed to demonstrate the automatic activation of color features, it turns out that the accessibility of these features varies across contexts. In the Simon task that has long been believed to demonstrate the automatic activation of spatial features, it turns out that the availability of these features varies across contexts as well. In general, context modulates the activation of many kinds of central features across diverse forms of conceptual knowledge. Given the ubiquitous effects of contexts on the central features of concepts, it appears increasingly likely that the activation of all conceptual knowledge is at least somewhat context-dependent.

For our purposes here, such findings suggest that the concepts combined during concept composition typically vary widely in content (also see Asher, 2011; Bosch, 1983). When a given concept is combined with other concepts, its content is likely to vary considerably across concept compositions. Even when the same two concepts are combined on different occasions, their individual content is likely to vary, especially when combined in different background situations.

Situation-specific conceptual content. A related phenomenon is that the varying content of a concept often appears to contain information relevant in the current situation. In some of these cases, the information appears to be stored in the concept, with the current situation selecting it as relevant (e.g., Barsalou, 1982; Conrad, 1978; Greenspan, 1986; Whitney et al., 1985). When processing the concept of *water* in the context of a Norwegian pond during winter, the property of *freezes* becomes active, although it would not become active normally (e.g., in the context of drinking *water* in a Norwegian restaurant). In other cases, varying content appears to originate in background situations (e.g., Barsalou & Wiemer-Hastings, 2005; Papies, 2013; Wu & Barsalou, 2009). When conceptualizing *chair*, for example, people might include chairs as occurring in offices. Similarly, when conceptualizing *truth*, people might represent truth as it occurs in courts of law. In general, considerable amounts of information in

conceptual representations appear to originate in the background situation in which a concept is currently represented, including the setting, related objects, other agents, actions, and variety of internal states (e.g., goals, evaluations, affect, motivation, reward, mentalizing).

Furthermore, the background situations active for a given category are likely to vary widely, both between and within individuals. The concept of *chair*, for example, is likely to be associated in memory with background situations for kitchens, living rooms, classrooms, theaters, jets, and ski lifts. Not only do different kinds of chairs (and chair properties) become active when representing chairs in these different background situations, so does differing knowledge about the external environment, action, and internal states. Sitting in a jet chair, for example, is associated with a different setting, actions, and internal states than sitting in a living room chair. Yeh and Barsalou (2006) review research in episodic and semantic memory showing that a given concept is often associated with a variety of background situations in memory. As a consequence, when a concept is processed in a particular situation, the situation activates related background information that enters into task performance.

Bayesian sampling. As we have just seen, the properties active for a concept vary with context. As we also saw earlier, concepts may not be associated with stable cores that are activated automatically. How, then, should we think about the varying content of a concept?

One possibility is that conceptual content is sampled in a Bayesian manner; in other words, sampling reflects both frequency of use and contextual relevance (Barsalou, 2011). In a given situation, information that has been processed frequently across situations for the concept has a higher probability of becoming active than information processed infrequently (e.g., round vs. floats for *basketball*). Importantly, however, high probability information need not be core information, becoming active on all occasions when the concept is processed. Rather than being active obligatorily across situations, this information simply has a higher probability of being active than less frequent information. Depending on the context, high probability information may or may not be active (Lebois et al., 2014).

Additionally, information that is relevant for the concept in the current situation is more likely to become active than information less relevant in the situation. Thus, floats has a higher probability of

becoming active for *basketball* in the context of needing a life preserver than in the context of playing pickup basketball. Putting these two factors together, the information that becomes sampled for a concept on a given occasion is a Bayesian function of its overall accessibility in memory and its situational relevance. Future research is needed to assess this conjecture, but given the general relevance of Bayesian processes in intelligent systems, Bayesian processes are likely to play central roles in conceptual processing as well.

General concepts. If the conceptual content represented for a concept varies dynamically across situations, how can a concept be represented in a general manner (cf. Pelletier, this volume). The grounded approach to concepts offers several potential solutions to this issue.

First, the distributed network in the brain that aggregates multimodal information for a concept across its exemplars offers a general representation of the respective category. Because the network aggregates a tremendous amount of information across experiences with the category, it doesn't simply represent a single exemplar, but represents all exemplars together (see the distinction between simulations and simulators in Barsalou, 1999, 2009; also Barsalou, in press-b). Although only one specific representation of the concept is constructed at a time dynamically from the network, the general concept is represented implicitly across the entire network. To the extent that people are aware that they possess much more knowledge about a concept than they can access in a given moment, they may believe that they possess a general representation at a meta-cognitive level. Furthermore, the entire network may enter implicitly and unconsciously into cognitive functions, thereby producing general effects of the concept.

Second, the most accessible simulation that can be easily constructed from the network offers a default representation of the concept (cf. McNally and Boleda, this volume). Although the most accessible simulation is a specific simulation (or perhaps a family of closely-related simulations), it nevertheless offers a prototypical representation of all the simulations that the network is capable of producing. As a result, it potentially functions as a general representation of all possible simulations, in the spirit of global attractors (cf. McClelland & Rumelhart, 1985) and prototype theory (cf. Hampton, 2006).

Third, people may learn to explicitly construct specific simulations of a concept that offer a

generic representation of it (e.g., one that omits situational detail). To the extent that people understand meta-cognitively that a concept can be general, they may be able to construct a simulation that captures this generic character. Notably, such simulations are simply some of the infinitely many simulations that the distributed multimodal network for a concept can construct. As a consequence, generic simulations reside together with situation-specific simulations that support situated action across diverse contexts. In a sense, generic simulations, too, are situation-specific simulations, with the relevant situations being those where generic representations are useful.

Sharing concepts. If the conceptual content represented for a concept varies dynamically, how can individuals ever come to effectively share a common conceptual representation in a specific situation (cf. Pelletier, this volume)? One likely solution to this problem follows from the facts that different individuals have similar bodies, brains, and cognitive systems; they live in similar physical environments; they operate in highly-coordinated social contexts. As a result, different individuals acquire similar distributed networks for a given concept over the course of development. Within a particular social group or culture, different individuals' networks are likely to be highly similar, given similar coordinated experiences with many shared exemplars. Even across different cultures, these networks are likely to be highly similar, given that all humans have similar bodies, brains, and cognitive systems, operating in similar physical and social environments.

As we have seen, however, the network for a concept is capable of producing an infinite number of dynamically-varying simulations across situations. How, then, do different individuals coordinate a shared simulation in a given situation, even when they share highly similar networks? One potential solution is that coordination results from establishing *common ground* (e.g., Clark, 1996). Through the use of linguistic reference, shared experience, cultural convention, and pragmatic strategies, people coordinate their use of a concept in the current situation, attempting to construct similar simulations of it from their similar networks. As a result, different individuals may typically be able to establish a shared simulation relevant in the current situation, such that effective understanding and communication follows. Although the simulations that different conversants construct are never identical, they may typically be similar enough for communicative success (with ineffective communication occurring when the

coordination of a shared simulation fails). Although much remains to be learned about how individuals coordinate dynamically constructed simulations, doing so seems feasible, given that establishing common ground and utilizing effective pragmatic strategies may effectively constrain the construction of simulations from shared conceptual networks.

Implications for concept composition. This section has reviewed three properties of concepts that present potential challenges for theories of concept composition. First, concepts typically do not appear to have stable cores. Second, the content active for a concept varies across situations. Third, the information active on a given situation for a concept may result from Bayesian sampling.

Together these three properties suggest that when a given concept is combined with others across situations, it is likely to exhibit no core stability and considerable situation-specific variability. Although representations of the concept across situations probably exhibit statistical regularities associated with Bayesian priors and situational relevance, they nevertheless take diverse forms, similar to how a given phoneme constantly adapts to its phonemic and articulatory contexts (e.g., Repp, 1982). To the extent that these three properties are present in concepts, they are likely to play significant roles in concept composition, thereby placing challenging constraints on accounts of this process.

Multiple Representational Forms

Exemplars vs. abstractions. Almost always, even psychological theories of concept composition assume that abstractions represent the concepts being combined, with these abstractions typically taking the form of rules, prototypes, and frames/schemata (e.g., Barsalou, 1992, 2012; Barsalou & Hale, 1993; Hampton, 2006; Murphy, 2002; Pothos & Wills, 2011). Nevertheless, considerable evidence has accumulated that a given concept includes detailed information about specific category exemplars (e.g., Allen & Brooks, 1991; Medin & Schaffer, 1978; Nosofsky, 2011; also see Minda & Smith, 2011). Rather than abstracting over the idiosyncratic features of category exemplars, conceptual knowledge includes these features and correlations between them. Furthermore, when people later use this knowledge, they often appear to draw on exemplar-level information. Such findings have suggested to many researchers that memories of a category's exemplars represent a concept, not an abstraction.

It is worth noting, however, that abstractions

(not just exemplars) are capable of representing detailed information about category exemplars (Barsalou, 1990, 2003a; McClelland & Rumelhart, 1985). Nevertheless, the important possibility remains that people often activate memories of specific exemplars when representing a concept. If so, then how does concept composition with the concept proceed? Even if an abstraction were to represent the concept, what are the implications for concept composition of exemplar-specific details that that the concept probably contains?

Amodal symbols. Typically, accounts of concept composition assume that it operates on representations that are amodal, symbolic, and propositional (e.g., Barsalou, 1999, 2008a, 2010). From this perspective, a symbolic representational structure, such as relation, predicate, frame, or schema, has arguments, variables, or slots, whose values vary across instances and situations (e.g., Barsalou, 1992). Given these kinds of representational structures, concept composition becomes readily feasible (e.g., Gagné & Spalding, 2014; Wisniewski, 1997). To implement the combination of an object with a property, for example, a variable in the frame for the object is set to a specific value. When representing the concept composition of *green flamingo*, for example, the *color* variable in the *flamingo* frame is set to the value *green*. Similarly, to implement the combination of two nouns, a frame assigns the two nouns as arguments of different variables. When representing the concept composition of *flamingo fare*, a frame for *eat* could take *flamingo* and *fare* as values of the variables for *agent* and *theme*, respectively.

Other accounts of concept composition assume that feature sets for individual concepts are combined in varying ways to produce the meanings of complex phrases (e.g., Hampton, 1997; Hampton & Jönsson, 2012). Again, the respective features are typically assumed to be amodal symbolic representations. When representing *sports that are games*, for example, people combine the feature sets for *sports* and *games* to produce the combined concept.

Amodal representations such as frames and features may play central roles in the representation of concepts, and in turn, during the process of combining them. Nevertheless, other important representational processes appear to play important roles in concept composition as well, in particular, multimodal simulations and linguistic forms. If so, then theories of concept composition may benefit from taking them into account.

Multimodal simulations. Earlier, in the

introduction, the construct of multimodal simulation was introduced when defining *concept*. To reiterate briefly, a distributed multimodal network for a category in the brain attempts to simulate the kind of neural and bodily states that occur while interacting with a category's members. When simulating a bicycle, for example, the brain and body reenact the kinds of states that occur while experiencing bicycles, including how they look, where they are used, how to ride one, and associated internal states. Considerable evidence has accumulated supporting the proposal that simulation is a basic computational process in the brain, not only in conceptual processing, but in all cognitive processes, ranging from perception to social cognition (e.g., Barsalou, 2008a, in press-b). Certainly, multimodal simulation is not the only representational process in the brain, but it appears to be a central one used widely and frequently. If so, then it may contribute significantly to the process of combining concepts.

To date, little research has addressed the role of simulation in concept composition directly. One set of experiments, however, provides some initial evidence (Wu & Barsalou, 2009). Across three experiments, a preliminary prediction was that when participants produced the features of an object, they would simulate an experience of the object, reporting features salient in the simulation. As a consequence, features occluded in the simulation, and therefore not salient, wouldn't be produced frequently. When participants produce features of *lawn*, for example, they should simulate a green lawn. Because *green* and *grass blades* are salient in the simulation, they should be produced frequently. Because *dirt* and *roots* are occluded, they should be produced rarely. Across experiments, salient features were produced much more often than occluded features, as the simulation account predicts.

The primary prediction across experiments was that when the same objects were combined with modifiers that revealed normally occluded features, these occluded features would become salient and therefore be reported more frequently. When participants received the concept composition of *rolled-up lawn*, for example, they should simulate a rolled-up lawn. Because *dirt* and *roots* become unoccluded in the simulation and therefore salient, they should be reported more often than for *lawn* alone. Across experiments, this finding occurred robustly. Not only did it occur for familiar phrases such as *half-watermelon*, it occurred for unfamiliar phrases such as *glass car*.

Most importantly, however, unoccluded

features did not become more salient for noun phrases such as *rolled-up snake*. This finding suggests two conclusions. First, modifiers such as *rolled-up* do not have symbolic rules associated with them specifying that the occluded features of any noun they modify become salient. If such rules existed, then normally occluded features should always become more salient whenever the modifier combines with any noun. Again, however, occluded features did not become more salient for phrases such as *rolled-up snake*. Second, modifiers such as *rolled-up* only change the salience of normally occluded features when the process of rolling up an object unoccludes these features, as when someone rolls up a lawn, but not when a snake rolls up itself. Because it's unlikely that representations of *rolled-up* with every possible noun are already stored in memory, it's likely that most interpretations of these noun phrases are computed online via concept composition, as the phrase is heard. Although it is in principle possible to develop an amodal procedure that correctly increases (or doesn't increase) the salience of occluded features across all *rolled up* noun phrases, a simpler approach exists: Simply simulate rolling up the noun object, and report the features salient in the process, especially in the end state. Rather than relying on a complex amodal rule to determine the concept composition, a simulation of the noun phrase naturally specifies whether occluded features become salient or not.

More generally, simulation may play central roles across many kinds of concept composition. Consider Medin and Shoben's (1988) example of *gray hair* vs. *gray cloud*, where *gray* takes on different values in the two noun phrases. Although amodal symbols could capture such values, a simpler and more intuitive account is that simulating *gray hair* and *gray cloud* lead to different shades of gray being simulated. Because people store memories of *gray hair* and *gray cloud*, they can easily retrieve and simulate these memories to produce meanings of the respective phrases. Similarly, consider the noun-noun combinations of *mountain trail* and *desert trail*. By retrieving and simulating memories of each, the resulting interpretations implicitly represent the features that *mountain trails* tend to change elevation and direction more often than do *desert trails*.

Finally, multimodal simulations capture affordances that may often occur during concept composition. The phrase *handleless cup*, for example, has the affordance of being grasped with a power grip rather than with a precision grip. On

hearing this phrase in a sentence like, "Would you enjoy drinking from a handleless cup," much work suggests that people would immediately produce a motor simulation of picking the cup up with a power grip to answer the question (e.g., Chaigneau, Barsalou, & Sloman, 2004; Heard, Masson, & Bub, 2015). Because multimodal simulations carry affordances of objects and situations, activating these simulations during the concept composition that underlies language use prepares comprehenders for effective situated action in the world.

As these examples illustrate, multimodal simulation offers a natural means of establishing the integration of conceptual information that occurs in concept composition. Not only does multimodal simulation offer a means of combining concepts never combined before as just described (e.g., *handleless cup*), it also offers a means of explaining the extensional feedback that underlies the learned representations of many familiar phrases (e.g., *tea cup*), as described later. Multimodal simulation is also compatible with the graphical diagrams used widely across cognitive linguistics to represent the concept composition that underlies spatial relations, part-whole relations, and event structure (e.g., Langacker, 1986, 1987, 2008; Talmy, 1983, 1988, 2000).

Distributed linguistic representations. Still another potentially important representational mechanism for implementing concept composition is the distribution of linguistic forms associated with a concept (with linguist forms including words, phrases, and sentences). Depending on the approach, a distributed linguistic representation can simply include the word associates of a concept, or it can be the surrounding words and syntax that typically occur in sentences that include the concept. Following Paivio's (1971, 1986) Dual Coding Theory, many cognitive psychologists have proposed that linguistic forms work together with an image-based conceptual system to produce intelligent computation across a wide variety of cognitive tasks (for several recent perspectives, see Barsalou, Santos, Simmons, & Wilson, 2008; Glaser, 1992; Louwerse, 2008, 2011; Louwerse & Connell, 2011).

Other researchers from the tradition of distributional semantics have similarly suggested that distributed linguistic representations play a wide variety of important roles in conceptual processing (Baroni & Lenci, 2010; Erk, 2012; Erk & Padó, 2008; Landauer & Dumais, 1997; Landauer, McNamara, Dennis, & Kintsch, 2013; Padó & Lapata, 2007). Characterizing the word for

a concept in terms of its distributed linguistic representation taps into the word's semantics, such that its similarity to the semantics of other words can be established. To the extent that two concepts tend to cooccur with the same linguistic forms, the more similar they are semantically. As much research shows, the distributed linguistic representation associated with a concept affects its processing across a wide variety of tasks and phenomena (e.g., Barsalou et al., 2008; Erk, 2012; Landauer et al., 2013; Louwerse, 2011). Thus, it wouldn't be surprising if distributed linguistic representations enter into concept composition as well. As McNally and Boleda (this volume) suggest, distributed linguistic representations could play central roles both in establishing both default interpretations of combined concepts and situationally-relevant interpretations (consistent with the Bayesian account of conceptual variability presented earlier).

Consider two further possibilities. First, the distributed linguistic representations of two words being combined may suggest relations for combining the underlying concepts. When establishing the meaning of *bicycle corridor*, for example, the distributed linguistic representations for *bicycle* and *corridor* may be established initially. Because *walk* is a salient word associate of *corridor*, it becomes active quickly. Once *walk* becomes active for *corridor*, it then matches the word for *walk* that is stored with *bicycle* but is currently inactive (because *ride* is more highly associated with *bicycle* than is *walk*). As a result of this match, *walk* becomes active for *bicycle*, leading to the interpretation of *bicycle corridor* as a place where one walks a bicycle. Clearly, other relational interpretations are possible, but this interpretation illustrates how combining distributed linguistic representations could help converge one of the many possible concept compositions available.

A second possibility is that the distributed linguistic representations of two combined words project into long-term memory together, activating memories that contain words from both. Imagine reading a newspaper article in the past about a bicycle race around the underground snow corridors of a Canadian university, with the phrase *bicycle corridor* not being mentioned. When hearing the phrase *bicycle corridor* on a later occasion, the distributed linguistic representations generated from the two words could map onto words from the earlier text, activate the associated memory, and influence how *bicycle corridor* is currently understood. Rather than being interpreted

as a corridor where bicycles are walked, *bicycle corridor* is now interpreted instead as a corridor where bicycles are raced.

Finally, researchers working from the distributed semantics perspective increasingly explore distributed representations of non-linguistic information, such as the images associated with a concept (e.g., Anderson, Bruni, Lopopolo, Poesio, & Baroni, in press; Thill, Padó, & Ziemke, 2014). To the extent that the distributions of images associated with two concepts are similar, their conceptual content is likely to be similar as well. Interestingly, the distribution of images associated with a concept may index the distributed multimodal networks that represent concepts in grounded approaches (cf. (Barsalou, 1999, 2009, in press-a, in press-b)).

Implications for concept composition. As this section has described, a concept can be associated with a wide variety of representational forms. A concept can be represented with abstractions, exemplars, or both. Orthogonally, these abstractions and exemplars could be represented with amodal symbols, multimodal simulations, distributed linguistic representations, or some combination. At this point, we don't really know which of these forms enter into concept composition. It wouldn't be surprising, though, if they all did, depending on the kinds of concepts being combined, together with the goals and constraints associated with a given processing context.

Pragmatic Constraint

Extensional feedback. Researchers argue increasingly that cognition cannot be studied independently of sensory-motor systems, the body, and the environment (e.g., Aydede & Robbins, 2009; Barsalou, in press, 2008a, 2010; Barsalou, Breazeal, & Smith, 2007; Clark, 1998, 2008). From this perspective, it may be impossible to understand concept composition adequately without taking into account how it is grounded in these important ways.

Hampton's (1988) construct of *extensional feedback* illustrates how the environment can contribute to concept composition (see McNally and Boleda, this volume, for the related construct of *referentially-afforded modification*). Knowing that a *pet bird* is usually caged probably doesn't result from conceptual knowledge stored with *pet* or with *bird*, given that neither is usually caged. Instead, this knowledge originates in seeing that the pet birds in one's experience typically live in cages. Thus, the meaning of *pet bird* is established by

activating memories of caged birds in the world—its meaning doesn't solely result from conceptually combining *pet* and *bird*.

The visual world paradigm developed by Tanenhaus, Spivey-Knowlton, Eberhard, and Sedivy, (1995) offers another example of extensional feedback (for a recent review, see Huettig, Rommers, & Meyer, 2011). In this paradigm, participants perceive a visual world and are requested to act on it. For example, a participant might see objects laid out on a table and be asked to move one object on top of another. As participants hear various phrases, such as “apple on a towel,” they have to figure out what these phrases mean. Much work demonstrates that the meanings of these phrases don't only result from concept composition; they also result from extensional feedback via the visual world. The visual world paradigm offers much potential for studying both concept composition and the role of extensional feedback.

Effects of extensional feedback are likely to vary across cultures, further demonstrating situational effects on concept composition. Imagine the classic example of a *pet fish*. For many westerners who have had goldfish and guppies as pet fish, the extensional feedback from these experiences probably contributes to how *pet fish* is interpreted. Imagine, though, a culture that doesn't have pet fish. What interpretations of *pet fish* would they produce? Without extensional feedback, they would have to rely on concept composition to construct interpretations, such that *pet fish* might well reflect a combination of what's typical for the two constituent concepts, *pet* and *fish* (e.g., Hampton, 1988; Hampton & Jönsson, 2012). Alternatively, imagine cultures that keep different kinds of pet fish than do westerners, for example, relatively wild fish kept in backyard ponds, or in walled-off areas of streams and rivers. For these cultures, extensional feedback from the respective experiences would produce still different interpretations of *pet fish*, again not based on concept composition but understood as a conventional phrase.

Multimodal simulation, discussed earlier, offers a natural mechanism for explaining how extensional feedback becomes established in memory. When experiencing a pet fish, for example, a multimodal sensory-motor representation of the fish is established, including its background situation. If someone uses the phrase “pet fish” to describe the fish in this situation, the phrase becomes stored with the multimodal memory. On hearing the phrase, “pet

fish,” on a later occasion, the associated memory becomes active as a simulation, which then represents the meaning of the phrase. In this way, the simple process of capturing and running multimodal simulations may contribute to the comprehension of concept compositions.

Finally, concept composition interacts with extensional feedback in perhaps unanticipated ways. Imagine never having heard the phrase, “pet fish” before and never having encountered a pet fish. On hearing someone say, “Meet my pet fish,” as they point to a goldfish in a glass tank, it is necessary to make sense of *pet fish* in this context. Even though an excellent example of a pet fish is readily perceptible in the situation, understanding the uttered phrase, *pet fish*, and mapping it onto its referent requires concept composition. Clearly, a variety of situational cues make this possible, but it is still probably necessary to combine *pet* and *fish* conceptually so that understanding what *pet fish* means can occur. This suggests that concept composition plays important roles in understanding immediate situations, and doesn't only occur during language comprehension about non-present situations. Counter-intuitively, perhaps, understanding a phrase or sentence about a non-present situation may often draw heavily on extensional feedback, whereas understanding what a phrase or sentence means in the current situation may sometimes rely more on concept composition.

Situational constraint. As described earlier, people's knowledge of a category includes knowledge of background situations in which its instances occur. Imagine hearing the sentence, “In my kitchen, I sat on an uncomfortable chair for three hours straight.” To comprehend the concept composition, *uncomfortable chair*, someone might represent the *chair* as a simple solid wooden chair of the type typically found in kitchens, and might represent *uncomfortable* as a sore bottom from sitting too long on hard wood. In contrast, imagine hearing, “On the jet, I sat in an uncomfortable chair for three hours straight.” Here someone might represent *chair* as a cramped jet chair, and *uncomfortable* as stiffness from not being able to move.

As these examples illustrate, the background situation constrains the meaning of a concept composition, such as *uncomfortable chair*, such that it takes on different meanings in different situations (for a review, see Yeh & Barsalou, 2006). With respect to the principle of variable content described earlier, situational constraint offers one central mechanism that causes conceptual content to change across situations. As

people comprehend words, phrases, and sentences, they heavily utilize information in the background situation to construct meaning, a point well-established in the comprehension literature (e.g., Bransford & Johnson, 1972, 1973). Although the principle of situational constraint is related to the earlier principle of situation-specific conceptual content, the two principles differ in emphasis. Whereas situation-specific conceptual content focuses on situational information as a source of varying conceptual content, situational constraint focuses on the pragmatic influence that the current situation has on establishing varying content.

Implications for concept composition. As this section has described, interactions with past and present situations appear to play central roles in how people combine concepts. Often, memories of experienced entities and events from past situations provide the meanings of current phrases and sentences via extensional feedback. Similarly, the relevant background situation constrains how concept composition proceeds, biasing the conceptual content activated and combined toward situationally relevant information. It is likely that many other forms of pragmatic constraint not anticipated here contribute to concept composition as well.

Conclusion

From the perspective of human cognition, principles of conceptual processing have potential implications for theories of concept composition. First, the content active for a concept on a given occasion appears to vary considerably. Rather than having a stable core, a concept exhibits dynamically varying content across situations, perhaps resulting from Bayesian sampling. Second, multiple representational forms may constitute the dynamically constructed representation of a concept on a given occasion. Although these representations could include abstractions, they may also include exemplars, and are at least likely to include exemplar-level information. Additionally, multimodal simulations and distributed linguistic representations may become active and play central roles in concept composition, along with amodal symbols. On a given occasion, the specific mix of representational forms may vary widely. Third, the dynamic representation of a concept on a given occasion is likely to reflect a variety of pragmatic constraints. On the one hand, these representations could contain conceptual content established through extensional feedback on earlier occasions. On the other, they could contain information associated

with the current background situation.

The principles of conceptual processing just described may be incorrect. Much controversy remains concerning the nature of the human conceptual system, such that some researchers, at least, would take issue with the principles of conceptual processing presented here. Nevertheless, based on the current empirical literatures in cognitive science, it is likely that at least some of these principles are approximately correct, and that the human conceptual system does not take the idealized forms assumed in some accounts of concept composition.

A possible framework for understanding concept composition. Let's assume for the moment that the account of conceptual processing presented here is roughly correct. For the sake of argument, how might we explain concept composition if these principles are true? What kinds of theories might be successful under conditions of varying conceptual content, multiple representational forms, and pragmatic constraint?

Under these conditions, it seems difficult to imagine how a formal symbolic theory could be successful. Perhaps I'm being uncharitable, but a successful theory of this type would require that considerable information about exemplars, situations, and multimodal content be formalized, as would be procedures for combining it (consider what would be required to handle the representations of *rolled-up lawn* vs. *rolled-up snake* described earlier).

Perhaps a more tractable approach is to assume that when people perform concept composition, they typically attempt to construct a multimodal simulation of a relevant situation (Barsalou, 1999, 2008b; Wu & Barsalou, 2009). In the process, people draw heavily on previous situational experience, as they attempt to construct a simulation that includes all the concepts being combined in a sensible manner. Certainly, various types of grammatical and linguistic patterns may contribute to this process. Nevertheless, attempting to construct a coherent simulation may be central.

Notably, simulation accounts of concept composition differ considerably from symbolic accounts. Whereas simulation accounts represent conceptual content implicitly in multimodal representations captured from situational experience, symbolic approaches abstract this information from situational experience and encoded it explicitly into amodal symbolic structures. As a consequence, the processes that underlie concept composition differ considerably at the level of cognitive and neural mechanisms (cf.

Marr, 1982). In simulation accounts, as we saw earlier for Wu and Barsalou (2009), complex simulations of experience are constructed and transformed, whereas in symbolic accounts, explicit symbolic structures are combined according to the rules of logic and language. Although, at an abstract level, the information being represented in the two approaches may be comparable, the representation and processing of this information differs considerably, leading to different explanations and predictions at the cognitive and neural levels.

Although it remains to be seen whether a simulation approach to concept composition will work, this approach has several qualities that might yield useful outcomes. First, the simulation approach relies heavily on situational experience, both past and current, consistent with the work reviewed earlier on the importance of situations for conceptual processing. Not only does the multimodal simulation constructed to represent a concept composition draw heavily on situational memories, it also draws pragmatically on the current physical situation.

Second, this approach seems well suited for handling variable conceptual content and multiple representational forms. From this perspective, it is not a problem that a concept's content varies constantly—it only matters whether its current content can be integrated with the evolving multimodal simulation of the relevant situation. Because the process of integrating simulations seems potentially “sloppy,” varying content might not constitute a severe problem. Similarly, it doesn't really matter whether an exemplar or an abstraction becomes active, because either could be integrated into an evolving simulation. Additionally, distributed linguistic representations could become active throughout the simulation process, activating and structuring relevant simulation content.

Third, this approach doesn't require a tight formalism to be effective. Rather than linguistic forms being associated with semantic formulae that combine in a compositional manner, linguistic forms and grammatical structures constitute instructions for constructing, transforming, and evaluating multimodal simulations (Barsalou, 1999; cf. Langacker, 1986, 1987, 2008). Because linguistic forms and grammar become cues for building simulations, their formal properties are less important than the effects that they have on the simulation process.

Clearly, however, many approaches for explaining concept composition exist, and many

new ones are likely to develop. Perhaps what is most significant and encouraging currently is that diverse groups of researchers increasingly attempt to understand concept composition from different perspectives (e.g., the contributors to the current volume and to Werning et al., 2012). Because concept composition may constitute an unusually challenging aspect of human cognition to explain, it seems likely that adopting different theoretical perspectives will be essential, as will be incorporating many kinds of empirical evidence.

Pelletier (this volume) offers a long list of challenges for cognitivist approaches to concepts, noting that cognitivist approaches have yet to provide accounts of singular terms, quantification, grammatical categories, syntax, integrated thoughts, negation, and so forth. Clearly, these are important issues, and cognitive approaches have done relatively little to address them. To a large extent this is because cognitive approaches have focused on other issues, often attempting to establish the most basic mechanisms in the brain that underlie conceptual processing, holding off on more challenging issues until basic issues are better understood. Additionally, cognitive researchers often lack the technical expertise required to perform informed empirical study of more complex issues.

Rather than challenging the cognitivists to explain these issues on their own, perhaps formalists could work together with cognitivists to develop accounts that integrate both perspectives. Ultimately, optimal accounts of conceptual processing are likely to reflect sophistication about both its formal and cognitive aspects.

Furthermore, to the extent that the formal aspects of conceptual processing reflect structure in the world (e.g., truth conditions), grounded cognition offers a natural theoretical framework for integrating the formal and cognitive perspectives. Because grounded cognition focuses on coupling cognition with the environment, it sees the structure of concepts in the world as indispensable to conceptual processing (although it might take a more statistical view of the physical world as opposed to a logical one). From the grounded perspective, a successful account of conceptual processing is only possible once both perspectives are included and unified. A pure cognitive account on its own is far from sufficient, as is a pure formal account of structure in the world. Because cognition is coupled with the world, taking each seriously in explaining the other is essential.

References

- Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, *120*, 3–19.
- Anderson, A., Bruni, E., Lopopolo, M., Poesio, M., & Baroni, M. (in press). Reading visually embodied meaning from the brain: Visually grounded computational models decode visual-object mental imagery induced by written text. *NeuroImage*.
- Asher, N. (2011). *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press.
- Augustinova, M., & Ferrand, L. (2014). Automaticity of word reading: Evidence from the semantic Stroop paradigm. *Current Directions in Psychological Science*, *23*, 343–348.
- Aydede, M., & Robbins, P. (2009). *The Cambridge handbook of situated cognition*. Cambridge: Cambridge University Press.
- Baroni, M., & Lenci, A. (2010). Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics*, *36*, 673–721.
- Barsalou, L. W. (in press-a). Can cognition be reduced to action? Processes that mediate stimuli and responses make human action possible. In A. K. Engel, K. J. Friston, & D. Kragic, *Where's the action? The pragmatic turn in cognitive science (Strüngmann Forum Reports, Vol. 18. J. Lupp, Series Ed.)*. Cambridge, MA: MIT Press.
- Barsalou, L. W. (in press-b). Situated conceptualization: Theory and application. In Y. Coello & M. H. Fischer, *Foundations of embodied cognition*. East Sussex: Psychology Press.
- Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, *10*, 82–93.
- Barsalou, L. W. (1989). Intraconcept similarity and its implications for interconcept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76–121). Cambridge: Cambridge University Press.
- Barsalou, L. W. (1990). On the indistinguishability of exemplar memory and abstraction in category representation. In T. K. Srull & R. S. W. Jr, *Content and process specificity in the effects of prior experiences: Advances in social cognition* (Vol. 3, pp. 61–88). Hillsdale, NJ: Erlbaum.
- Barsalou, L. W. (1992). Frames, concepts, and conceptual fields. In A. Lehrer & E. F. Kittay (Eds.), *Frames, fields, and contrasts: New essays in semantic and lexical organization* (pp. 21–74). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–660.
- Barsalou, L. W. (2003a). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London. Series B: 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. (2008a). Grounded Cognition. *Annual Review of Psychology*, *59*, 617–645.
- Barsalou, L. W. (2008b). Grounding symbolic operations in the brain's modal systems. In G. R. Semin & E. R. Smith, *Embodied grounding: Social, cognitive, affective, and neuroscientific approaches* (pp. 9–42). New York: Cambridge University Press.
- Barsalou, L. W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*, 1281–1289.
- Barsalou, L. W. (2010). Grounded cognition: past, present, and future. *Topics in Cognitive Science*, *2*, 716–724.
- Barsalou, L. W. (2011). Integrating Bayesian analysis and mechanistic theories in grounded cognition. *Behavioral and Brain Sciences*, *34*, 191–192.
- Barsalou, L. W. (2012). The human conceptual system. In M. Spivey, K. McRae, & M. F. Joannisse, *The Cambridge handbook of psycholinguistics* (pp. 239–258). New York: Cambridge University Press.
- Barsalou, L. W., Breazeal, C., & Smith, L. B. (2007). Cognition as coordinated non-cognition. *Cognitive Processing*, *8*, 79–91.
- Barsalou, L. W., & Hale, C. (1993). Components of conceptual representation. From feature lists to recursive frames. In I. Van Mechelen, J. A. Hampton, R. Michalski, & P. Theuns, *Categories and concepts: Theoretical views and inductive data analysis* (pp. 97–144). San Diego: Academic Press.
- Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg, & A. C. Graesser, *Symbols, embodiment, and meaning* (pp. 245–283). Oxford: Oxford University Press.
- Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. A. Zwaan, (pp. 129–163). New York: Cambridge University Press.
- Barsalou, L. W., Yeh, W., Luka, B. J., Olseth, K. L., Mix, K. S., & Wu, L.-L. (1993). Concepts and meaning. In K. Beals, G. Cooke, D. Kathman, K. E. McCulloch, S. Kita, & D. Teste, *Chicago Linguistics Society 29: Papers from the parasession on conceptual representations* (pp. 23–61). University of Chicago: Chicago Linguistics Society.
- Blanchette, I., & Dunbar, K. (2000). How analogies are generated: The roles of structural and superficial similarity. *Memory & Cognition*, *28*, 108–24.
- Bosch, P. (1983). Vagueness" is context-dependence. A solution to the Sorites paradox. In T. T. Ballmer & M.

- Pinkal, *Approaching vagueness* (pp. 189–210). Amsterdam: North Holland.
- Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, *11*, 717–726.
- Bransford, J. D., & Johnson, M. K. (1973). Considerations of some problems of comprehension. In W. G. Chase, *Visual information processing* (pp. 383–438). Oxford: Academic Press.
- Casasanto, D., & Lupyan, G. (in press). All concepts are ad hoc concepts. In E. Margolis & S. Laurence, *The conceptual mind: New directions in the study of concepts*. Cambridge, MA: MIT Press.
- Chaigneau, S. E., Barsalou, L. W., & Sloman, S. A. (2004). Assessing the causal structure of function. *Journal of Experimental Psychology: General*, *133*, 601–625.
- Clark, A. (1998). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Connell, L., & Lynott, D. (2014). Principles of representation: Why you can't represent the same concept twice. *Topics in Cognitive Science*, *6*, 390–406.
- Conrad, C. (1978). Some factors involved in the recognition of words. In J. W. Cotton & R. L. Klatzky, *Semantic factors in cognition* (pp. 103–120). Hillsdale, NJ: Erlbaum.
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: constraint-guided conceptual combination. *Cognitive Science*, *24*, 299–349.
- Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain*. New York: WW Norton & Company.
- Donald, M. (1993). Precursor of Origins of the modern mind: Three stages in the evolution of culture and cognition. *Behavioral and Brain Sciences*, *16*, 737–748.
- Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition*, *110*, 412–431.
- Erk, K. (2012). Vector space models of word meaning and phrase meaning: A survey. *Language and Linguistics Compass*, *6*, 635–653.
- Erk, K., & Padó, S. (2008). A structured vector space model for word meaning in context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 897–906). Stroudsburg, PA, USA: Association for Computational Linguistics. Retrieved from
- Evans, V. (2009). *How words mean: Lexical concepts, cognitive models, and meaning construction*. Oxford: Oxford University Press.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*, 141–205.
- Gagné, C. L., & Spalding, T. L. (2014). Conceptual composition: The role of relational competition in the comprehension of modifier-noun phrases and noun-noun compounds. *The Psychology of Learning and Motivation*, *59*, 97–130.
- Gawronski, B., & Cesario, J. (2013). Of mice and men: What animal research can tell us about context effects on automatic responses in humans. *Personality and Social Psychology Review*, *17*, 187–215.
- Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*, 524–575.
- Glaser, W. R. (1992). Picture naming. *Cognition*, *42*, 61–105.
- Greenspan, S. L. (1986). Semantic flexibility and referential specificity of concrete nouns. *Journal of Memory and Language*, *25*, 539–557.
- Hampton, J. A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 12.
- Hampton, J. A. (1997). Conceptual combination. In K. Lamberts & D. R. Shanks, *Knowledge, concepts, and categories* (pp. 133–159). East Sussex: Psychology Press.
- Hampton, J. A. (2006). Concepts as prototypes. In B. H. Ross, *The psychology of learning and motivation: Advances in research and theory (Vol 46)* (pp. 79–113). San Diego, CA, US: Elsevier Academic Press.
- Hampton, J. A., & Jönsson, M. (2012). Typicality and compositionality: The logic of combining vague concepts. In M. Werning, W. Hinzen, & E. Machery, *The Oxford handbook of compositionality* (pp. 385–482). Oxford: Oxford University Press.
- Heard, A. W., Masson, M. E., & Bub, D. N. (2015). Time course of action representations evoked during sentence comprehension. *Acta Psychologica*, *156*, 98–105.
- Huetting, F., Rommers, J., & Meyer, A. S. (2011). Using the visual world paradigm to study language processing: A review and critical evaluation. *Acta Psychologica*, *137*, 151–171.
- Jones, M., & Love, B. C. (2011). Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, *34*, 169–188.
- Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, *23*, 221–236.
- Kiefer, M., Adams, S. C., & Zovko, M. (2012). Attentional sensitization of unconscious visual processing: Top-down influences on masked priming. *Advances in Cognitive Psychology*, *8*, 50–61.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to

- Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.
- Landauer, T. K., McNamara, D. S., Dennis, S., & Kintsch, W. (2013). *Handbook of Latent Semantic Analysis*. East Sussex: Psychology Press.
- Langacker, R. W. (1986). An introduction to cognitive grammar. *Cognitive Science*, *10*, 1–40.
- Langacker, R. W. (1987). *Foundations of Cognitive Grammar: Theoretical prerequisites*. Stanford: Stanford University Press.
- Langacker, R. W. (2008). *Cognitive grammar: A basic introduction*. Oxford: Oxford University Press.
- Lebois, L. A. M., Wilson-Mendenhall, C. D., & Barsalou, L. W. (2014). Are automatic conceptual cores the Gold Standard of semantic processing? The context-dependence of spatial meaning in grounded congruency effects. *Cognitive Science*, in press.
- Louwerse, M. M. (2008). Embodied relations are encoded in language. *Psychonomic Bulletin & Review*, *15*, 838–844.
- Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, *3*, 273–302.
- Louwerse, M. M., & Connell, L. (2011). A taste of words: Linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, *35*, 381–398.
- Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition*, *104*, 19–46.
- Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology-Paris*, *102*, 59–70.
- Malt, B. C. (1994). Water is not H₂O. *Cognitive Psychology*, *27*, 41–70.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY: Henry Holt.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*, 159–188.
- McCloskey, M., & Glucksberg, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology*, *11*, 1–37.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, *20*, 158–190.
- Minda, J. P., & Smith, J. D. (2011). Prototype models of categorization: Basic formulation, prediction, and limitations. In E. M. Pothos & A. J. Wills, *Formal approaches in categorization* (pp. 40–64). New York: Cambridge University Press.
- Murphy, G. L. (1988). Comprehending Complex Concepts. *Cognitive Science*, *12*, 529–562.
- Murphy, G. L. (2002). *The big book of concepts*. MIT Press.
- Padó, S., & Lapata, M. (2007). Dependency-based construction of semantic space models. *Computational Linguistics*, *33*, 161–199.
- Paivio, A. (1971). *Imagery and verbal processes*. Oxford: Holt, Rinehart & Winston. Retrieved from
- Paivio, A. (1986). *Mental representations: A dual-coding approach*. Oxford: Oxford University Press.
- Papies, E. K. (2013). Tempting food words activate eating simulations. *Frontiers in Psychology*, *4*, 1–12.
- Pothos, E. M., & Wills, A. J. (2011). *Formal approaches in categorization*. Cambridge: Cambridge University Press.
- Repp, B. H. (1982). Phonetic trading relations and context effects: New experimental evidence for a speech mode of perception. *Psychological Bulletin*, *92*, 81–110.
- Ross, B. H. (1984). Reminders and their effects in learning a cognitive skill. *Cognitive Psychology*, *16*, 371–416.
- Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 629–639.
- Simmons, W. K., & Barsalou, L. W. (2003). The Similarity-in-Topography Principle: Reconciling Theories of Conceptual Deficits. *Cognitive Neuropsychology*, *20*, 451–486.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, *81*, 214.
- Talmy, L. (1983). How language structures space. In H. Pick & L. Acredelo, *Spatial orientation: Theory, research, and application* (pp. 225–282). New York: Plenum Press.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, *12*, 49–100.
- Talmy, L. (2000). *Toward a cognitive semantics, Vol. 1: Concept structuring systems*. Cambridge, MA: The MIT Press.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632–1634.
- Thill, S., Padó, S., & Ziemke, T. (2014). On the Importance of a Rich Embodiment in the Grounding of Concepts: Perspectives From Embodied Cognitive Science and Computational Linguistics. *Topics in Cognitive Science*, *6*, 545–558.
- Tomasello, M. (2009). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.

- Werning, M., Hinzen, W., & Machery, E. (2012). *The Oxford handbook of compositionality*. Oxford: Oxford University Press.
- Whitney, P., McKay, T., Kellas, G., & Emerson, W. A. (1985). Semantic activation of noun concepts in context. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 126–135.
- Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, *4*, 167–183.
- Wu, L. L., & Barsalou, L. W. (2009). Perceptual simulation in conceptual combination: Evidence from property generation. *Acta Psychologica*, *132*, 173–189.
- Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *The American Journal of Psychology*, *119*, 349–384.