Acta Psychologica 132 (2009) 173–189

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Perceptual simulation in conceptual combination: Evidence from property generation

Ling-ling Wu [a], Lawrence W. Barsalou [b,*]

[a] *Department of Information Management, National Taiwan University, Taiwan*
[b] *Department of Psychology, Emory University, Atlanta, GA 30322, USA*

ARTICLE INFO

Available online 18 March 2009

PsycINFO classification:
2340

Keywords:
Perceptual simulation
Conceptual combination
Property generation
Feature listing
Situated cognition
Concepts and categories
Knowledge
Mental imagery

ABSTRACT

In three experiments, participants received nouns or noun phrases for objects and verbally generated their properties ("feature listing"). Several sources of evidence indicated that participants constructed perceptual simulations to generate properties for the noun phrases during conceptual combination. First, the production of object properties for noun phrases depended on occlusion, with unoccluded properties being generated more often than occluded properties. Because a perceptual variable affected conceptual combination, perceptual simulations appeared central to combining the concepts for modifiers and head nouns. Second, neutral participants produced the same distributions of properties as participants instructed to describe images, suggesting that the conceptual representations used by neutral participants were similar to the mental images used by imagery participants. Furthermore, the property distributions for neutral and imagery participants differed from those for participants instructed to produce word associations. Third, participants produced large amounts of information about background situations associated with the object cues, suggesting that the simulations used to generate properties were situated. The experiments ruled out alternative explanations that simulation effects occur only for familiar noun phrases associated with perceptual memories and that rules associated with modifiers produce occlusion effects. A process model of the property generation task grounded in simulation mechanisms is presented. The possibility of integrating the simulation account of conceptual combination with traditional accounts and well-established findings is explored.

1. Introduction

Conceptual combination is a fundamental process in human cognition. Besides knowing thousands of individual concepts, people can combine these concepts into an infinite number of more complex concepts. Conceptual combination is found throughout higher cognition. During the perception of a novel scene, concepts that categorize scene components are combined to interpret the novel configuration of entities and events perceived. During language production, speakers combine concepts as they conceptualize what to convey in an utterance. In turn, listeners combine concepts for words in sentences to conceptualize what speakers are saying. During thought, agents combine concepts as they solve problems, reason, and make decisions. Although non-humans may have some ability to combine concepts (e.g., when perceiving novel scenes), humans appear qualitatively superior.

In a classic paper, Fodor and Pylyshyn (1988) argued that conceptual combination is a signature ability of humans. Using a finite number of concepts, humans can construct an infinite number of new concepts productively whose meanings are systematically related to one another. Although major arguments exist about the mechanisms responsible for this ability—as described next—there is no doubt that humans have a powerful ability to combine concepts.

According to traditional theories of cognition, conceptual combination results from combining amodal symbols in predicate-calculus-like representations (e.g., Barsalou, 1992; Fodor, 1975; Fodor & Pylyshyn, 1988; Pylyshyn, 1984; Smith, Osherson, Rips, & Keane, 1988). Consider combining concepts for *blue* and *house* to represent the meaning of the noun phrase, "blue house."[1] According to traditional theories, *house* is a predicate-like structure (often referred to as a "frame" or "schema") that contains many arguments, each taking multiple values (e.g., arguments such as *color, size, and location*). To represent the meaning of "blue house," a symbol for the value *blue* becomes bound to the *color* argument in the predicate structure for *house*, creating the expression, *house (color = blue, size = ?, location = ?, ...)*. Thus, standard theories assume that

* Corresponding author. Tel.: +404 727 4338.
E-mail address: barsalou@emory.edu (L.W. Barsalou).
URL: http://userwww.service.emory.edu/~barsalou/ (L.W. Barsalou).

[1] Throughout this article, we use quotes to indicate the word or phrase for a concept (e.g., "house"), and italics to indicate the respective concept (e.g., *house*).

0001-6918/$ - see front matter
doi:10.1016/j.actpsy.2009.02.002

combining amodal symbols in a psychological predicate calculus underlies conceptual combination, and indeed much of cognition.

Alternative theories exist as well. In response to Fodor and Pylyshyn's classic (1988) article, connectionists developed statistical approaches for implementing conceptual combination (e.g., Pollack, 1990; Smolensky, 1990). Although these approaches embraced the functional importance of conceptual combination in human cognition, they argued that statistical mechanisms implement this functionality instead of predicate-calculus-like mechanisms. We, too, assume that statistical structures underlie conceptual combination, and that conceptual combination is a highly dynamical process (e.g., Barsalou, 2003a, 2005a; Barsalou, Breazeal, & Smith, 2007).

Still another alternative—the one pursued here—is that simulation mechanisms in the brain play a central role in implementing conceptual combination (Barsalou, 1999, 2003a; Barsalou, 2005a, 2008b). According to this account, humans represent individual concepts with simulations of their referents, which could take the form of exemplars or prototypes. When representing *house*, for example, people represent houses with multi-modal simulations of what the experience of a house is like, including perceptions of houses, actions taken on houses, and introspections about houses (i.e., emotions and beliefs). To perform conceptual combination, simulations of individual concepts are combined to form larger, more complex simulations (e.g., combining simulations of *blue* and *house* to simulate *blue house*). Notably, the simulation account does not assume that holistic bit-map-like images represent concepts. Instead, this account assumes that the construct of a simulator enables the symbolic functionality of classic predicate-calculus-like theories. Essentially, simulators function as the concepts in classic theories, making it possible to implement argument binding, conceptual combination, propositional interpretation, and inference. This account further assumes that the construction of simulations from simulators is highly dynamical, much in the spirit of connectionist approaches. Thus, the simulation account embraces all three perspectives: It implements classic predicate-calculus-like functionality using simulation mechanisms that are inherently statistical.

2. Paradigm

Three experiments assessed the simulation account of conceptual combination using the property generation task (also known as "feature listing"). For decades, property generation has been an important tool for measuring conceptual representations (e.g., Hampton, 1979; McRae, Cree, Seidenberg, & McNorgan, 2005; Rosch & Mervis, 1975; Rosch, Mervis, Johnson, Gray, & Boyes-Braem, 1976; Smith, 1978). Typically, researchers assume that property generation provides a window on the underlying representation of a concept, with the properties that participants generate reflecting the concept's content in long-term memory. Researchers across diverse areas, including cognitive psychology, social psychology, consumer psychology, clinical psychology, and neuropsychology, continue to use property generation to measure conceptual content of interest in their respective domains of study.

On each trial in the experiments here, participants received a noun or noun phrase for an object and generated typical properties for it. Thus, a participant might have received the noun "lawn" and generated the properties *plant, green, soft, blades, you play on it,* and *makes me feel good*. The protocol that a participant produced was parsed into a sequence of individual properties, which were then classified using a hierarchical coding scheme (Appendix A). At the general level in this scheme, each property was coded as a taxonomic category, an entity property, a situational property, or an introspective property. For example, when *lawn* was the concept, *plant* was coded as a taxonomic category, *blades* as an entity property, *you play on it* as a situational property, and *makes me feel good* as an introspective property. Within these four general categories, each property was assigned to 1 of 37 hierarchically nested specific categories. Thus, *plant* was coded as a *superordinate* category, *blades* as an *external component* of an entity, *you play on it* as situational *action*, and *makes me feel good* as an *evaluative* introspection. For other applications of this coding scheme, see Cree and McRae (2003), McRae et al. (2005), and Brainerd, Yang, Reyna, Howe, and Mills (2008).[2]

In the experiments to follow, application of this coding scheme to property generation data was used to test three hypotheses associated with the simulation account of conceptual combination: (1) Occlusion modulates the generation of properties from nouns and noun phrases, (2) Participants spontaneously use simulations to generate properties, such that the properties they generate are roughly equivalent to those that participants instructed explicitly to use imagery generate (i.e., instructional equivalence), (3) Participants situate the simulations that represent nouns and noun phrases in the context of settings, events, and mental states. The next three sections address these hypotheses in turn.

3. Occlusion modulates property generation during conceptual combination

Much behavioral research demonstrates that when participants generate mental images in working memory, perceptual variables such as orientation, position, and size affect processing (e.g., Finke, 1989; Kosslyn, 1980; Shepard & Cooper, 1982). Because imagery utilizes representations in perceptual systems, variables that affect perception also affect imagery. Just as it takes longer to scan across 30° of visual angle in the visual field than 10°, it also takes longer to scan across 30° of a visual image than 10°. The following experiments pursue a similar research strategy: If participants construct perceptual simulations spontaneously to represent concepts, then perceptual variables should affect conceptual processing, analogous to how they affect perception.

Occlusion was the perceptual variable manipulated in the experiments here. In vision, occluded properties are not visible, whereas unoccluded properties are. As a result, occluded properties are less available for processing than unoccluded properties. Analogously, if participants simulate the referents of concepts to generate properties, then occluded properties in these simulations should be more difficult to generate than unoccluded properties. To see this, imagine generating properties for *lawn*. If participants simulate a normal experience of a lawn, *roots* and *dirt* should be relatively difficult to generate because they are occluded beneath the lawn's surface. Because *roots* and *dirt* are difficult to perceive for a lawn in actual perception, they should also be difficult to perceive in a simulation. Conversely, imagine generating properties for *rolled-up lawn*. If participants now simulate a lawn that has been rolled up, *roots* and *dirt* are exposed in the simulation and should be generated more readily. Because *roots* and *dirt* are easy to perceive for a rolled-up lawn in actual perception, they should also be easy to perceive in a simulation.

The experiments here used conceptual combination to test these predictions. In each experiment, some participants generated properties for isolated nouns (e.g., *lawn*). If these participants use simulations to represent these nouns, then occluded properties in these simulations should be generated at low rates (e.g., *roots and dirt*). Other participants in each experiment generated proper-

[2] The coding scheme in Appendix A has evolved since its use in the experiments here. For recent versions of this coding scheme, contact the second author and the authors of the papers just cited.

ties for conceptual combinations (e.g., *rolled-up lawn*). If these participants combine simulations for the modifier and head noun to construct the noun phrase's meaning, then occlusion should again affect the properties generated. Because, however, these noun phrases contain revealing modifiers (e.g., *rolled-up*), properties that were occluded for the isolated nouns should now become unoccluded for the noun phrases. Whereas *roots* and *dirt* are occluded for *lawn,* they become unoccluded when a simulation of *lawn* is combined with a simulation of *rolled-up*. If the simulation account of conceptual combination is correct, occlusion should modulate the generation of internal properties, depending on whether a revealing modifier is present or absent.

Traditional accounts of conceptual combination do not predict a priori that revealing modifiers should affect property generation in this manner. According to these accounts, the amodal symbol for *blades* does not literally occlude the amodal symbol for *roots* in the concept for *lawn*. *Blades* and *roots* may reside at different levels of accessibility in memory, with *blades* being more accessible than *roots*, but *blades* does not literally occlude *roots*. Furthermore, when "rolled-up" modifies "lawn," *roots* does not become unoccluded, given that it was never occluded in the first place. As this example illustrates, traditional accounts do not predict a priori that revealing modifiers should increase the production of occluded properties. Because occlusion is only a property of perceptual representations, only perceptual accounts of conceptual processing predict a priori that revealing modifiers should increase the accessibility of occluded properties.

The simulation account makes a further prediction about the effects of revealing modifiers during conceptual combination: When a revealing modifier makes internal properties visible, it typically does so at the expense of external properties, which become less visible. Again compare simulations of *lawn* vs. *rolled-up lawn*. In a simulation of *lawn*, the external properties of a lawn, such as *green* and *blades*, cover much of the simulation. In a simulation of *rolled-up lawn*, however, the amount of external lawn visible typically decreases, as more of the lawn's underside, and the surface below it, comes into view. Thus, we further predicted that revealing modifiers would not only increase the number of internal properties generated for a noun phrase, but that they would also simultaneously decrease the number of external properties. In other words, the simulation account predicts that there should be an inverse relationship between the production of internal and external properties across nouns and noun phrases.

Because traditional theories have unlimited expressive power—by virtue of being grounded in predicate calculus—it may well be possible for them to explain this complex pattern of predicted occlusion effects (e.g., Anderson, 1978). Notably, however, amodal theories typically do not make such predictions a priori. Furthermore, it is important that these theories explain the effects of perceptual variables in a principled manner, not from an unconstrained post hoc perspective (Barsalou, 1999, 2008a).

4. Instructional equivalence

The simulation account predicts that participants should spontaneously represent concepts with simulations without any instruction to do so. Assessments of instructional equivalence in the experiments here tested this hypothesis.

4.1. Imagery versus neutral instructions

In one condition of the experiments to follow—the neutral condition—participants received standard instructions for property generation, simply being asked to generate the characteristic properties of each concept, with nothing being said about imagery. In a second condition—the imagery condition—participants were asked to generate properties from mental images. When receiving the noun "lawn," for example, they were asked to construct an image of a lawn and then to describe the image's properties. Much research demonstrates that participants construct images when asked to do so (e.g., Finke, 1989; Jeannerod, 1995; Kosslyn, 1980; Shepard & Cooper, 1982).

Of interest is the similarity of the properties that neutral participants generate to the properties that imagery participants generated. According to the simulation account, neutral participants should have spontaneously represented the concepts with simulations of their instances. For example, when neutral participants generated properties for *lawn*, they should have simulated a lawn, scanned across it, and reported the properties they perceived (see Barsalou, 2003a; Barsalou, 2005a, for a more detailed account of how the simulation account explains the process of property generation). Most importantly, if neutral participants simulated instances to generate properties, their performance should have been comparable to the performance of imagery participants. When detailed analysis of the properties generated by both conditions is performed, we should find high correlations between the distributions of properties that neutral and imagery participants generated—what we will refer to as "instructional equivalence."[3]

The simulation account does *not* predict that imagery and neutral participants should produce identical distributions of properties. According to this account, the schematic framework underlying a simulation can be instantiated with many perceptual properties or just a few (e.g., Barsalou, 1999, 2003a; Barsalou, 2005a, 2008b). Thus, simulations can range from detailed and complete to schematic and incomplete. Furthermore, highly detailed simulations may tend to be relatively conscious, whereas more schematic simulations may tend to be relatively unconscious. In the experiments here, we therefore predicted that imagery participants would construct richer simulations than neutral participants. Nevertheless, because both groups use simulations to generate properties, their simulations should still exhibit strong similarities.

In contrast, traditional theories do not predict instructional equivalence a priori. To the contrary, many predict instructional inequivalence. Typically, traditional theories assume that the mechanisms underlying perception and action differ significantly from the mechanisms underlying conception. Thus, these theories predict that imagery participants should use simulations to generate properties, whereas neutral participants should use predicate-calculus-like structures. Because imagery and neutral participants use different mechanisms to generate properties, their performance should differ. Participants in the two conditions should produce different distributions of properties that reflect fundamentally different types of representation. If traditional theories are correct, instructional *inequivalence* should be observed for the neutral and imagery conditions.

4.2. Word association instructions

The experiments to follow included an additional production mode—the word condition. Specifically, this third group of participants was asked to produce words associated to the word for each

[3] We assume that high similarity between two distributions of properties should typically reflect correlations of approximately .80 and higher (uncorrected for reliability). Most importantly, however, we predict that correlations should be *relatively* higher when participants in both the neutral and imagery conditions produce properties from the simulation system than when neutral participants produce properties from simulations and word participants produce properties from the linguistic system (as described shortly).

target concept. For example, these participants were asked to generate words associated with the word "lawn."

The word condition extends the predictions of instructional equivalence and inequivalence as follows. Traditional theories, besides assuming that amodal symbols represent concepts, often assume that concepts correspond closely to words (e.g., Collins & Loftus, 1975; Fodor, 1975; Pylyshyn, 1984; Smith, 1978). In English, for example, the amodal symbol that represents the concept *lawn* corresponds to the word "lawn." If one further assumes that the organization of concepts parallels the organization of words, at least approximately, instructional equivalence between the neutral and word conditions follows. Because the organization of concepts roughly mirrors the organization of words, it should not matter whether neutral participants access concepts to generate responses or whether word participants access words. The distributions of properties that both groups generate should be similar.

In contrast, the simulation account predicts instructional inequivalence between the neutral and word conditions. According to the simulation account, neutral participants should construct simulations to represent concepts, whereas word participants should primarily produce lexical associates of the target word, bypassing the conceptual system (Barsalou, Santos, Simmons, & Wilson, 2008; Glaser, 1992). As a result, significant differences should occur between the neutral and word conditions.

Thus, the simulation account and traditional theories make different predictions about instructional equivalence and instructional inequivalence, which are summarized as follows:

Simulation account: imagery instructions $\approx$ neutral instructions $\neq$ word instructions
Traditional theories: imagery instructions $\neq$ neutral instructions $\approx$ word instructions

Whereas the simulation account predicts approximate equivalence between imagery and neutral instructions, it predicts inequivalence between neutral and word instructions. Conversely, traditional theories often predict approximate instructional equivalence between neutral and word instructions, but predict inequivalence between neutral and imagery instructions.

5. Situated simulations

When attention focuses on an object in perception, the background situation is always present. The situation does not disappear but recedes into the background. If participants simulate perceptual experience to represent concepts, then the concepts that they simulate should analogously be situated in appropriate contexts (Barsalou, 2003b, 2005b, 2008c; Barsalou, Niedenthal, Barbey, & Ruppert, 2003; Yeh & Barsalou, 2006). Even though participants are asked explicitly to produce object properties, they should also produce properties of the background situations that arise implicitly while simulating objects. When asked to generate properties of *lawn*, for example, participants might simulate a lawn in a park during festival, not just an isolated patch of grass. In other words, they simulate "being there" in a situation that contains the target entity (Barsalou, 2002). Furthermore, in the process of generating properties, participants not only generate properties of the target entity but of the background situation as well. Properties of the background situation leak into participants' protocols.

To test this hypothesis, the following experiments assessed whether participants produced situational properties that lie outside the target objects. Not only were external situations, such as a lawn or park of interest, so were participants' mental states about a target object's relevance. For example, participants might state that *lawn* has the property of *beautiful*, expressing a mental evaluation of it. Thus, situations were defined as both the external and mental contexts of the target entities.

6. Overview of the experiments

Experiment 1 assessed all three hypotheses of interest. First, it assessed the effects of occlusion, namely, do revealing modifiers increase the production of internal properties while reducing the production of external properties? Second, does instructional equivalence occur between the neutral and imagery conditions but not between the neutral and word conditions? Third, are the simulations of concepts that participants construct situated, such that participants not only produce properties of the target concepts but also properties of the surrounding situations?

Experiment 2 addressed the possibility that simulation mechanisms construct novel conceptual combinations productively (Barsalou, 1999, 2003a, 2005a, 2008b). It could be argued that the occlusion effects observed for the conceptual combinations in Experiment 1 only occurred because people have perceptual memories of these combinations from experiencing their referents (e.g., *convertible car*). Because these combinations are associated with perceptual memories, occlusion affects the generation of properties from them. This argument further proposes that the meanings of novel combinations—not associated with perceptual memories by definition—can only be constructed by combining predicate-calculus-like structures. Because non-perceptual representations underlie the construction of these novel meanings, generating their properties should not produce occlusion effects.

Conversely, the simulation account proposes that people combine simulations of familiar concepts to construct representations of novel concepts. To represent the novel concept of *glass car*, for example, people simulate a car and then change the simulation of its body from metal to glass. If perceptual simulations underlie the meanings of novel combinations, then occlusion effects should occur for them as well as for familiar combinations. To test this hypothesis, Experiment 2 assessed whether occlusion effects occur for both familiar and novel combinations. Observing occlusion effects for both supports the view that simulation mechanisms construct novel combinations productively. Observing occlusion effects only for familiar combinations supports the view that novel combinations require predicate-calculus-like structures.

Finally, Experiment 3 explored whether rules associated with revealing modifiers underlie the increased accessibility of occluded properties. If such rules exist, they should increase the accessibility of occluded properties whenever a revealing modifier is combined with a noun. For example, if *rolled-up* is associated with a rule that makes internal properties more accessible, then the accessibility of internal properties should increase not only for *rolled-up lawn* but also for *rolled-up snake*. To test this hypothesis, Experiment 3 assessed whether revealing modifiers ever fail to increase the accessibility of internal properties. If such failures occur, this would suggest that rules do not underlie occlusion effects and that simulations do instead.

7. Experiment 1

Three groups of participants generated properties under different production modes: imagery, neutral, or word association instructions. Crossed orthogonally with production mode, half of the participants in each production mode generated properties for one of the two concept types: nouns or noun phrases, where the noun phrases contained the nouns preceded by the revealing modifiers. For all analyses that followed, individual properties in the protocols were coded hierarchically at 4 general levels and at 37 specific levels (Appendix A).

To test for the effects of occlusion, proportions of internal properties and external properties (identified through the coding process) were computed for each concept for each participant. Of interest was whether the proportion of internal properties increased for noun phrases relative to nouns, and whether the proportion of external properties decreased.

To test for instructional equivalence, the overall proportion for each of the 37 specific property types in Appendix A was computed for each of the 6 conditions in the production mode X concept type design. Each vector of 37 proportions constituted a profile of the property information generated by the respective condition. Of interest was whether the correlations between vectors for the neutral and imagery conditions were high, and whether the correlations between the neutral and word conditions were low.

To test whether participants used situated representations to produce concepts, the overall proportion for each of the 4 general property types in Appendix A was computed for each condition. Of interest was whether significant amounts of non-entity properties were produced, especially properties that described situations and introspections.

7.1. Method

7.1.1. Design and participants

Two independent variables crossed orthogonally between participants structured the experiment: production mode (imagery, neutral, or word instructions) and concept type (nouns or noun phrases). Participants were 24 volunteers at the University of Chicago, with 4 assigned randomly to each cell in the production mode X concept type design. Each participant generated properties for 10 concepts, received in one of 3 random orders, distributed equally across conditions.

7.1.2. Materials

Table B1 in Appendix B presents the 10 object concepts in Experiment 1. As Table B1 illustrates, the noun phrase condition used the same 10 nouns as the noun condition, each preceded by a revealing modifier. Each noun referred to an entity whose internal properties are normally occluded by an external surface (e.g., *lawn*), whereas each noun phrase referred to an entity whose internal properties become visible with the modifier (e.g., *rolled-up lawn*). Three random orders of the nouns were constructed, with the same three orders used for the corresponding noun phrases.

7.1.3. Procedure

Participants performed the experiment individually. They first read instructions that described the production mode they would be using. The imagery instructions emphasized that the task tested people's ability to describe images. These participants were explicitly asked to construct an image for each noun or noun phrase and then to describe the contents of the image. The neutral instructions described the stimuli as "concepts" and emphasized that participants' task was to generate characteristics typically true of each one. These instructions were as neutral as possible so as not to bias participants towards imagery or word association. The word association instructions emphasized that participants' task was to produce words associated with each noun or a noun phrase. Participants were asked to produce whatever words came to mind immediately on hearing each item. In all conditions, participants were told that they would receive around a dozen items, and that their spoken protocols would be tape recorded. On each trial, participants heard one of the following probes:

- *Imagery*: construct an image of a ______ and describe the contents of your image.
- *Neutral*: describe the characteristics typically true of the concept ______.
- *Word Association*: freely generate whatever words immediately come to mind for ______.

Participants received either the 10 nouns or the 10 nouns phrases sequentially, with each noun or noun phrase inserted into the blank slot of the relevant probe.

7.2. Results

7.2.1. Coding the protocols

After transcribing a participant's protocol for a concept, it was parsed into individual properties. The coding scheme in Appendix A guided this process. Because this scheme captures detailed property information about concepts, individual words in a protocol, not just phrases, were often coded. As described earlier, each property was coded at a general level (taxonomic category, entity property, situational property, introspective property, miscellaneous category), and also at a specific level (1 of 37 nested codes). Appendix A provides examples of coded items from the protocols. Two judges coded every protocol at the specific level and achieved agreement of 91%. Disagreements were resolved via subsequent discussion.

7.2.2. Statistical analysis

In computing inferential tests for all experiments, both participants and concepts were treated as random factors in F' ratios (Clark, 1973). Because F' ratios take both participants and concepts into consideration, they are relatively conservative. When assessing a priori predictions, we applied the Dunn-Sidak procedure for non-orthogonal contrasts. These planned comparisons will be denoted as $t_{DS}(c,v)$, where c is the number of comparisons, and v is the degrees of freedom in the error term for the variable tested (Kirk, 1982). When multiple comparisons between means are performed, the Dunn-Sidak procedure is also relatively conservative, providing protection against Type I errors.

7.2.3. Overall performance

Participants produced 8.56 properties per concept on average. More properties were produced for noun phrases (10.16) than for nouns (6.96), with this difference approaching significance ($F'(1, 19) = 3.24$, $MS_e = 5.18$, $p < .08$). Production mode affected property frequency significantly ($F'(2,20) = 7.38$, $MS_e = 5.33$, $p < .01$), with imagery participants (12.04) producing more properties than neutral participants (9.84), who produced more properties than word participants (3.79) ($t_{DS}(2,20) = 5.75$, $MS_e = 5.33$, $p < .01$; $t_{DS}(2, 20) = 16.71$, $MS_e = 5.33$, $p < .001$, respectively). Concept type and production mode did not interact ($F'(2,18) = 1.12$, $MS_e = 5.05$).

7.2.4. Occlusion

The simulation account predicts that occlusion should affect the proportion of internal and external properties that participants produce. When internal properties are occluded in the noun condition, their accessibility should be lower than when they are unoccluded in the noun phrase condition. Conversely, as internal properties become more accessible in the noun phrase condition, external properties should become less accessible.

For the purpose of this analysis (and the analogous analyses in Experiments 2 and 3), internal properties were defined as both internal surface properties and internal component properties, whereas external properties were defined as both external surface properties and external component properties (Appendix A). The proportion of internal properties was computed for each protocol that a participant produced for a concept, as was the proportion of external properties. Each proportion was the frequency of

Table 1
Proportions of internal and external properties for nouns and nouns phrases (NPs) in Experiment 1.

	Imagery		Neutral		Word	
Property	Nouns	NPs	Nouns	NPs	Nouns	NPs
Internal	.11	.27	.17	.26	.20	.32
External	.71	.33	.51	.30	.44	.30

Table 2
Correlations between production modes across specific property types in Experiment 1.

Concept type	Correlated production modes		
	Neutral and imagery	Neutral and word association	Imagery and word association
Nouns	.89	.73	.70
Noun phrases	.96	.29	.30

internal or external properties divided by the total frequency of all entity properties (Appendix A). All proportions were submitted to an arcsin transformation to stabilize variance before submitting them to inferential tests (Winer, 1971), with MS_e values for these tests reported in arcsin units. Table 1 presents the average proportions from this analysis, prior to the arcsin transformation. Proportions in Experiments 2 and 3 were analyzed analogously.

Planned comparisons assessed the predictions that the proportion of internal properties should be higher for noun phrases than for nouns, whereas the proportion of external properties should be higher for nouns than for noun phrases. The following contrast—applied individually to each protocol—was used to test these combined predictions: internal properties for noun phrases (+1), internal properties for nouns (-1), external properties for noun phrases (-1), and external properties for nouns (+1). In support of the simulation account, this contrast across participants and protocols was significant both in the imagery condition ($t_{DS}(4,126) = 5.22$, $MS_e = .53$, $p < .001$) and in the neutral condition ($t_{DS}(4,126) = 2.94$, $MS_e = .53$, $p < .02$), and was marginally significant in the word condition ($t_{DS}(4,126) = 2.46$, $MS_e = .53$, $p < .06$). As predicted, occlusion modulated the production of internal and external properties in an inversely related manner.

Specific contrasts within these interactions further supported predictions of the simulation account. Overall, the proportion of internal properties was higher for noun phrases (.28) than for nouns (.16). This difference occurred for every production mode, being significant in the imagery condition ($t_{DS}(4,126) = 2.74$, $MS_e = .53$, $p < .03$), and being marginally significant both in the neutral condition ($t_{DS}(4, 126) = 1.95$, $MS_e = .53$, $p < .10$) and in the word condition ($t_{DS}(4, 126) = 2.04$, $MS_e = .53$, $p < .07$). Conversely, the overall proportion of external properties was higher for nouns (.55) than for noun phrases (.31). Again, this difference occurred for every production mode, being significant both in the imagery condition ($t_{DS}(4,126) = 4.65$, $MS_e = .53$, $p < .001$) and in the neutral condition ($t_{DS}(4,126) = 2.21$, $MS_e = .53$, $p < .05$), and being in the predicted direction but not significant in the word condition ($t_{DS}(4,126) = 1.43$, $MS_e = .53$).

7.2.5. Instructional equivalence

The simulation account predicts that both neutral and imagery participants should construct simulations to generate properties. If so, then neutral and imagery participants should produce similar distributions of properties. In contrast, traditional theories predict that neutral and imagery participants use different cognitive systems to produce conceptual information. If so, then different distributions of properties should be observed for neutral and imagery participants.

To test these predictions, the average frequency for each of the 37 property types was computed across participants and concepts for each of the 6 conditions in the production mode X concept type design. The resulting vector of 37 average frequencies for a condition constitutes a performance profile, characterizing the information that these participants produced. To assess instructional equivalence, these profiles were then correlated for each pair of conditions. Table 2 presents these correlations.

The results support the simulation account. The profiles for neutral and imagery participants correlated .93, averaged across nouns and noun phrases. In contrast, the profiles for neutral and word participants only correlated .51. Thus, neutral participants were much more similar to imagery participants than to word participants, suggesting that neutral participants used simulation to generate properties like imagery participants.

Two further analyses support this conclusion. First, the correlations in Table 2 were corrected for attenuation due to the unreliability of their average frequencies (Guilford & Fruchter, 1973). Although these averages were reasonably reliable (.82), correlations computed across them were attenuated, because their reliabilities were less than 1. To estimate the population values of the correlations in Table 2, the reliabilities of the average frequencies were entered into the standard correction for attenuation (Guilford & Fruchter, 1973). The corrected correlations between neutral and imagery participants increased from .93 to .99, indicating a very high similarity between these two groups in the information that they produced. In contrast, the corrected correlations between the neutral and word participants only increased to .64.

The second analysis used partial correlations to establish how much unique variance in the neutral condition could be explained by variance in the imagery and word conditions. Removing the variance of word participants from the correlations between neutral and imagery participants produced a partial correlation of .87 for properties. Because this value is nearly identical to the original (.93), word association played relatively little role in the neutral condition, whereas simulation of the sort used by imagery participants dominated. Conversely, removing the variance of imagery participants from the correlations between neutral and word participants produced a partial correlation of .17. This large decrease further indicates the importance of simulation in the neutral condition.

7.2.6. Situated simulation

Analysis of general property types assessed the prediction that that the simulations used to represent nouns and noun phrases were situated. If this hypothesis is correct, then participants should have generated many non-entity properties that described situations and introspections. The results in Table 3 confirm this prediction.

Because participants were asked explicitly to produce entity properties, it is not surprising that they produced entity properties more frequently than any other property type ($F'(3, 60) = 42.75$, $MS_e = 2.48$, $p < .001$). Entity properties (5.31) were more frequent than situational properties (1.78), introspective properties (.51), and taxonomic categories (.38), which did not differ. The surprising finding was how many non-entity properties participants produced, given that these properties were not cued. Consistent with the hypothesis that conceptualizations of objects are situated, 26% of the properties that participants produced described situations and introspections. These data strongly suggest that participants situated their object conceptualizations with respect to physical settings and mental perspectives.

Table 3
Proportions and frequencies of general property types in Experiment 1.

Property	Proportion	Frequency					
		Nouns			Noun phrases		
		Image	Neutral	Word	Image	Neutral	Word
Taxonomic	.04	.55	.53	.50	.15	.23	.30
Entity	.62	7.73	4.45	1.83	8.48	7.88	1.53
Situation	.20	1.00	1.65	1.38	2.70	2.58	1.38
Introspection	.06	.50	.18	.03	1.18	.88	.33
Total		10.10	6.80	3.95	13.98	12.88	3.63

Note: Because miscellaneous properties are not included above, the frequencies of category, entity, situation, and introspection properties do not sum to the total number of properties, nor do their proportions sum to 1.

7.3. Discussion

The results of Experiment 1 support the a priori predictions of the simulation account. Occlusion had substantial effects on property generation. When nouns referred to objects whose internal properties were occluded, these properties were relatively inaccessible. When noun phrases referred to the same objects with their interiors revealed, the accessibility of internal properties increased substantially. Furthermore, the increased accessibility of internal properties for noun phrases was accompanied by a decrease in the accessibility of external properties. As the internal properties of an object became more salient in a simulation, external properties became less salient.

Interestingly, occlusion effects occurred for every production mode, including word association. This finding suggests that information from simulations becomes available automatically and quickly during conceptual processing, even when the focus is on words. For results that further support this conclusion, see Pulvermüller, Shtyrov, and Ilmoniemi (2005), Solomon and Barsalou (2004), and Santos, Simmons, Chaigneau, and Barsalou (2008). Because occlusion effects were weaker for word association, however, and because instructional equivalence did not occur between word association and the other two conditions, additional mechanisms were probably responsible for generating properties. As suggested by Barsalou et al. (2008), mixtures of word association and simulation may often occur during conceptual processing. If so, the word condition here may have mostly used word association to generate properties, plus a smaller amount of simulation. As a result, this condition exhibited occlusion effects but not instructional equivalence.

The presence of instructional equivalence between the neutral and imagery conditions further supports the simulation account. As this account predicts, neutral and imagery subjects produced highly similar distributions of properties. After being corrected for attenuation, the correlation was .99, indicating that both groups produced essentially the same information. In contrast, neutral and word participants produced different property distributions, given that the correlations between them were much lower. Finally, partial correlations indicated that imagery performance explained nearly all the variance in the neutral condition, and that word performance explained hardly any. This overall pattern of results indicates that neutral participants used simulation much like imagery subjects.

The presence of situational and introspective information in participants' protocols demonstrates that the simulations used to generate properties were situated. Specifically, 26% of the properties generated did not describe target objects per se but instead described their situational and introspective contexts. This finding suggests that participants did not simply simulate the target objects in isolation but instead simulated them in background situations (Barsalou, 2003b, 2005b, 2008c; Yeh & Barsalou, 2006). Once these situated simulations became active, participants described all aspects of them, including entities, setting, events, and introspections.

Finally, imagery participants produced more properties than did neutral participants (12.04 vs. 9.84 per concept). As suggested earlier, imagery instructions are likely to induce conscious imagery that is relatively rich in detail, whereas neutral instructions are likely to induce imagery that is more unconscious and sketchier in detail. The difference in property frequency between imagery and neutral participants is consistent with this proposal. Interestingly, however, both groups showed strong occlusion effects, and exhibited a high degree of instructional equivalence. Although the amount of perceptual detail in simulations appeared to vary between groups, much underlying perceptual structure appeared to remain constant.

8. Experiment 2

Based on the results of Experiment 1, traditional theories might grant that simulations play a role in representing some conceptual combinations. Perhaps when people perceive the referents of familiar combinations, such as *half watermelon,* they store perceptual memories that later represent these combinations in simulations. From seeing half watermelons on picnics, for example, people store perceptual memories that they use later to simulate the meaning of "half watermelon." Because perceptual memories of "half watermelon" exist, there is no need to combine the concepts of *half* and *watermelon* to construct its meaning. Conceptual combination is not necessary because the retrieval of perceptual memories suffices.

Furthermore, theorists often argue that perceptual representations cannot be combined symbolically to form novel conceptual representations productively and systematically (e.g., Fodor, 1975). Instead conceptual combination can only be performed using predicate-calculus-like representations. If so, then traditional theories predict that people must use predicate-calculus-like representations to represent novel conceptual combinations. Because people have never seen referents of these combinations, and because they cannot combine perceptual memories for individual concepts symbolically to represent them, predicate-calculus-like representations must be used.

Experiment 2 addressed this hypothesis by manipulating whether participants generated properties for familiar noun phrases (e.g., *convertible car*) or for novel noun phases (e.g., *glass car*). As in Experiment 1, the effects of occlusion and instructional equivalence were examined. If participants use perceptual memories to represent familiar combinations, but use predicate-calculus-like representations to represent novel combinations, then occlusion effects and instructional equivalence should only occur for the familiar noun phrases. Alternatively, if occlusion effects and instructional equivalence occur for both familiar and novel combinations, this would support the proposal that simulation mechanisms implement conceptual combination (Barsalou, 1999, 2003a; Barsalou, 2005a, 2008b). Because unfamiliar noun phrases are not associated with perceptual memories, their meanings must be established by combining simulations for individual concepts.

8.1. Method

8.1.1. Design and participants

Two independent variables crossed orthogonally between participants structured the experiment: production mode (imagery, neutral, or word instructions) and concept type (familiar noun phrases, novel noun phrases, or nouns). Participants were 72 students at the University of Chicago paid for participating, with 8 assigned randomly to each cell in the production mode X concept type design. None had participated in a previous experiment. Four additional participants were replaced because they produced property frequencies that were two standard deviations above or below the mean for at least four critical concepts. Each participant generated properties for 10 target concepts and for 10 filler concepts, received in one of 3 random orders, distributed equally across conditions.

8.1.2. Materials

Table B2 in Appendix B presents the critical materials in Experiment 2. Each of the 10 nouns referred to an entity whose internal properties are normally occluded by an external surface (e.g., *car*). In the familiar and novel noun phrase conditions, a revealing modifier preceded each noun. In the familiar noun phrase condition, the 10 noun phrases referred to familiar objects whose interiors are exposed (e.g., *convertible car*). In the novel noun phrase condition, the 10 noun phrases referred to novel objects whose interiors are also exposed (e.g., *glass car*).

The critical materials in Table B2 were drawn from a preliminary study that scaled 23 familiar noun phrases and 20 novel noun phrases, which shared a common set of nouns. One group of 16 participants rated these 43 noun phrases for "How often have you heard this phrase," and a second group of 16 participants rated them for "How often have you encountered this object." Each group used a scale of never (1), rarely (2), occasionally (3), often (4), and very often (5). Participants in both groups were students at the University of Chicago who participated for pay and who had not participated in a previous experiment. Based on these ratings, the 10 most novel noun phrases were selected whose combined average ratings on phrase familiarity (1.24) and referent familiarity were the lowest (1.54), approaching the "never" scale value. Conversely, the 10 most familiar noun phrases were selected whose combined average ratings were the highest (2.69, 2.92), approaching the "occasionally" scale value. Table B2 presents the values for each noun phrase.

The critical materials were then scaled to ensure that the sizes of the exposed interiors were comparable for familiar and novel noun phrases. A third group of 12 participants received each critical noun first, followed by its familiar noun phrase and its novel noun phrase in a random order. For each of these three concepts, participants were asked to image its referent, and then to estimate what percentage of its interior was exposed. To achieve psychophysical calibration, participants first saw four drawings of an apple with its interior exposed 0%, 5%, 50%, or 95% and were told these percentages. Participants then received the 10 critical sets of materials with 8 filler sets in one of two random orders. Participants were students at the University of Chicago who participated for pay and who had not participated in a previous experiment. On average, participants reported that 63.61% of the 10 familiar objects were exposed and that 67.22% of the 10 novel objects were exposed. Because these percentages did not differ ($F(1,9) = .80$, $MS_e = 976.21$, *ns*), size of the exposed area was not confounded with familiarity.

Ten additional object concepts served as fillers to mask the presence of revealing modifiers in the list. In the noun conditions, the filler concepts were unmodified (e.g., *shoes* and *door*). In the noun phrase conditions, these concepts were modified but not with revealing modifiers (e.g., *expensive shoes and squeaky door*). The same filler nouns appeared across all conditions, preceded by the same filler modifiers in all noun phrase conditions. Three random orders of the 20 nouns were constructed, with the same three orders used for the corresponding noun phrases.

8.1.3. Procedure

The procedure was the same as in Experiment 1.

8.2. Results

The 10 critical protocols of 11 randomly selected participants were coded by two independent raters. Because agreement was 91%, only one judge coded the remainder of the protocols.

8.2.1. Overall performance

Participants produced 13.28 properties per concept on average. Number of properties did not differ between nouns (14.16), novel noun phrases (13.76), and familiar noun phrases (11.91) ($F'(2,63) = .63$, $MS_e = 14.91$). Production mode, however, affected property frequency ($F'(2,66) = 29.06$, $MS_e = 15.47$, $p < .001$), with imagery participants (20.37) producing more properties than neutral participants (15.35), who produced more properties than word participants (4.11) ($t_{DS}(2, 66) = 13.98$, $MS_e = 15.47$, $p < .001$; $t_{DS}(2,66) = 31.30$, $MS_e = 15.47$, $p < .001$). Concept type and production mode did not interact ($F'(4,62) = 0.68$, $MS_e = 14.74$).

8.2.2. Occlusion

The simulation account predicts that occlusion should affect the representation of both novel and familiar noun phrases. For both, internal properties should be relatively inaccessible for nouns when they are occluded but should become more accessible for noun phrases when they are unoccluded. Conversely, as internal properties became more accessible for noun phrases, external properties should become less accessible. Although the mechanism underlying these effects may differ for novel noun phrases (conceptual combination) and familiar novel noun phrases (referent retrieval), participants should construct simulations to represent both, thereby producing occlusion effects.

As in Experiment 1, the proportion of internal properties was computed for each protocol that a participant produced for a concept, as was the proportion of external properties. Table 4 presents the average proportions from this analysis. Also as in Experiment 1,

Table 4
Proportions of internal and external properties for nouns, novel noun phrases (Nov-NPs), and familiar noun phrases (Fam-NPs) in Experiment 2.

Property	Imagery			Neutral			Word association		
	Nouns	Nov-NPs	Fam-NPs	Nouns	Nov-NPs	Fam-NPs	Nouns	Nov-NPs	Fam-NPs
Internal	.11	.29	.39	.17	.30	.40	.26	.34	.41
External	.56	.35	.31	.34	.20	.22	.47	.28	.22

planned comparisons assessed the predictions that the proportion of internal properties should be higher for noun phrases than for nouns, whereas the proportion of external properties should be higher for nouns than for noun phrases. For each production condition, one contrast was computed for novel noun phrases relative to nouns, and a second contrast was computed for familiar novel noun phrases relative to nouns. Again one contrast was applied to each protocol using the following weights: internal properties for noun phrases (+1), internal properties for nouns (-1), external properties for noun phrases (-1), and external properties for nouns (+1).

In support of the simulation account, these contrasts were significant across participants and protocols. Most importantly, all three contrasts for the novel noun phrases relative to the nouns were significant: imagery condition ($t_{DS}(4,396) = 5.58$, $MS_e = .56$, $p < .001$), neutral condition ($t_{DS}(4,396) = 4.01$, $MS_e = .56$, $p < .001$), word condition ($t_{DS}(4,396) = 3.55$, $MS_e = .56$, $p < .001$). All three contrasts were also significant for the familiar noun phrases relative to the nouns: imagery condition ($t_{DS}(4,396) = 7.52$, $MS_e = .56$, $p < .001$), neutral condition ($t_{DS}(4,396) = 4.61$, $MS_e = .56$, $p < .001$), and word condition ($t_{DS}(4,396) = 5.20$, $MS_e = .56$, $p < .001$). As predicted, occlusion modulated the production of internal and external properties in an inversely related manner for both types of noun phrases.

Specific contrasts within these interactions further supported predictions of the simulation account. Overall, the proportion of internal properties was higher for novel noun phrases (.31) than for nouns (.18). This difference was significant both in the imagery condition ($t_{DS}(2,396) = 4.12$, $MS_e = .56$, $p < .01$) and in the neutral condition ($t_{DS}(2,396) = 2.63$, $MS_e = .56$, $p < .01$), and was marginally significant in the word condition ($t_{DS}(2,396) = 1.97$, $MS_e = .56$, $p < .10$). Overall, the proportion of internal properties was also higher for familiar noun phrases (.40) than for nouns (.18). This difference was significant in all three conditions: imagery ($t_{DS}(2,396) = 5.92$, $MS_e = .56$, $p < .001$), neutral ($t_{DS}(2,396) = 4.18$, $MS_e = .56$, $p < .001$), and word ($t_{DS}(2,396) = 3.23$, $MS_e = .56$, $p < .01$).

Conversely, the overall proportion of external properties was higher for nouns (.46) than for novel noun phrases (.28). This difference was significant in all three conditions: imagery ($t_{DS}(2,396) = 3.77$, $MS_e = .56$, $p < .01$), neutral ($t_{DS}(2,396) = 3.05$, $MS_e = .56$, $p < .01$), and word ($t_{DS}(2,396) = 3.05$, $MS_e = .56$, $p < .01$). The overall proportion of external properties was also higher for nouns (.46) than for familiar noun phrases (.25). This difference was significant in all three conditions: imagery ($t_{DS}(2,396) = 4.72$, $MS_e = .56$, $p < .001$), neutral ($t_{DS}(2,396) = 2.33$, $MS_e = .56$, $p < .05$), and word ($t_{DS}(2,396) = 4.12$, $MS_e = .56$, $p < .01$).

8.2.3. Instructional equivalence

As in Experiment 1, average frequencies across the 37 property types were computed to construct profiles for the information generated in each of the 9 conditions of the production mode X concept type design. Table 5 presents the correlations computed across these profiles to assess instructional equivalence. As in Experiment 1, these correlations support predictions of the simulation account. The profiles for neutral and imagery participants correlated .88, averaged across nouns, novel noun phrases, and familiar noun phrases. In contrast, the profiles for neutral and word participants only correlated .61. As in Experiment 1, neutral participants were more similar to imagery participants than to word participants, suggesting that neutral participants used simulations to generate properties like imagery participants.

Two further analyses further support the conclusion that neutral participants constructed simulations. First, the correlations in Table 5 were corrected for attenuation due to the unreliability of their average frequencies. Although these averages were reasonably reliable (.89), the corrected correlation between neutral and imagery participants increased from .88 to .96, indicating a very high similarity between these two groups in the information that they produced. In contrast, the corrected correlations between the neutral and word participants only increased to .70.

The second analysis used partial correlations to establish how much unique variance in the neutral condition could be explained by the variance in the imagery and word conditions. Removing the variance of word participants from the correlations between neutral and imagery participants produced a partial correlation of .79. Because this value approaches the original (.88), word association played relatively little role in the neutral condition, whereas perceptual simulation of the sort used by imagery participants dominated. Conversely, removing the variance of imagery participants from the correlations between neutral and word participants yielded a partial correlation of .20. These low values further indicate the importance of simulation in the neutral condition.

8.2.4. Situated simulation

As Table 6 illustrates, participants produced more entity properties than any other type ($F'(3,60) = 12.52$, $MS_e = 6.22$, $p < .001$). Entity properties (6.37) were more frequent than situational properties (3.59), which were more frequent than introspective properties (1.06), which were more frequent than taxonomic categories (.40) ($t_{DS}(3,60) = 21.15$, $MS_e = 6.22$, $p < .00001$; $t_{DS}(3,60) = 19.25$, $MS_e = 6.22$, $p < .001$; $t_{DS}(3, 60) = 5.02$, $MS_e = 6.22$, $p < .01$, respectively). Again, this is not surprising given that participants were explicitly instructed to produce entity properties. Of greater interest is the finding that participants situated their representations of nouns and noun phrases. Consistent with the hypothesis that simulations of objects are situated, 35% of the properties described situations and introspections. As in Experiment 1, these data strongly suggest that participants situated their object simulations in settings, events, and mental perspectives.

8.3. Discussion

The results of Experiment 2 further indicate that simulation underlies conceptual combination. Most significantly, occlusion effects and instructional equivalence occurred for novel noun phrases. If representing the meanings of novel noun phrases could not be performed using simulations—requiring predicate-calculus-like structures instead—occlusion effects would not have occurred. Their presence indicates that simulation mechanisms were operating during conceptual combination.

Again, occlusion effects occurred not only in the imagery and neutral conditions, but also in the word condition, suggesting that simulation mechanisms became active during word association. Because these effects were weaker, however, and because instructional equivalence did not occur between word association and the other two conditions, other processes appeared involved in generating properties. Again, the word condition may use a mixture of word association and simulation to generate properties (Barsalou et al., 2008).

Table 5
Correlations between production modes across specific property types in Experiment 2.

Concept type	Correlated production modes		
	Neutral and imagery	Neutral and word assoc.	Imagery and word assoc.
Nouns	.88	.59	.64
Novel noun phrases	.87	.50	.46
Familiar noun phrases	.88	.74	.73

Table 6
Proportions and frequencies of general property types in Experiment 2.

Property	Proportion	Frequency								
		Nouns			Novel noun phrases			Familiar noun phrases		
		Image	Neutral	Word	Image	Neutral	Word	Image	Neutral	Word
Taxonomic	.03	1.03	.80	.36	.74	.14	.13	.18	.33	.18
Entity	.48	12.16	7.46	2.01	1.18	5.96	1.45	1.13	6.48	1.94
Situation	.27	4.76	5.34	1.95	5.66	4.23	1.38	3.24	4.10	1.34
Introspection	.08	1.05	1.24	.20	2.43	2.21	.19	1.03	1.48	.29
Total		21.09	16.49	4.91	23.41	14.56	3.31	16.61	15.01	4.11

Note: Because miscellaneous properties are not included above, the frequencies of category, entity, situation, and introspection properties do not sum to the total number of properties, nor do their proportions sum to 1.

The presence of extensive situational and introspective information in participants' protocols—a total of 35%—again demonstrates that the simulations used to generate properties were situated. Once these situated simulations became active, participants described many aspects of them, including entities, setting, events, and introspections.

Finally, imagery participants produced more properties than did neutral participants (20.37 vs. 15.35). Imagery instructions again appeared to induce simulations that were more conscious and detailed than those in the neutral condition. Nevertheless, the presence of strong occlusion effects for both groups and a high degree of instructional equivalence indicates much commonality in how they generated properties.

9. Experiment 3

Traditional theories could explain occlusion effects as the result of rules that increase the importance of internal properties in object representations. To see this, consider the modifier *rolled-up*. Perhaps a rule stored with *rolled-up* specifies that the internal properties of the nouns it modifies should increase in accessibility. Thus, when *rolled-up* combines with *lawn*, the accessibility of *dirt* and *roots* increases, not because they become unoccluded in a simulation, but because a rule increases their accessibility in a predicate-calculus-like structure.

If such rules exist, then a revealing modifier should generally increase the accessibility of a noun's internal properties, regardless of the noun modified. To test this hypothesis, Experiment 3 assessed whether revealing modifiers fail to increase the accessibility of a noun's internal properties. For example, does *rolled-up* ever fail to increase the accessibility of internal properties in a noun it modifies, such as when it modifies *snake* in *rolled-up snake*?

The simulation theory does not predict that revealing modifiers should generally increase the accessibility of a noun's internal properties. When combined with some nouns, these modifiers do not refer to objects whose interiors are revealed. For example, the interior of a snake is not exposed in *rolled-up snake*. Whereas rules predict occlusion effects whenever a revealing modifier modifies a noun, the simulation account only predicts occlusion effects for those objects whose interiors become revealed. In the Discussion, we consider accounts that assign an individual occlusion rule to each noun phrase. As we will see, simulations motivate these rules and therefore explain occlusion effects more parsimoniously and more naturally.

To assess whether rules or simulations underlie the increased accessibility of internal properties for revealing modifiers, each revealing modifier was paired with two different nouns. For example, *rolled-up* was paired with *lawn* and *snake*. Whereas Experiments 1 and 2 held nouns constant and varied modifiers, Experiment 3 held modifiers constant and varied nouns. One group of participants received revealing noun phrases (e.g., *rolled-up lawn*), whereas a second group received non-revealing noun phrases (e.g., *rolled-up snake*). A third group received the nouns from the internal noun phrases (e.g., *lawn*), and a fourth group received the nouns from the non-internal noun phrases (e.g., *snake*). These latter two groups provided baseline levels of internal properties. If participants use rules to construct concepts for the noun phrases, then internal properties should increase above the noun baselines for both revealing and non-revealing noun phrases. Alternatively, if participants use simulations, internal properties should only increase for revealing noun phrases.

Because the two previous experiments exhibited consistent results for instructional equivalence, we did not manipulate production mode in Experiment 3. Instead, all participants received neutral instructions.

9.1. Method

9.1.1. Design and participants

Two independent variables crossed orthogonally between participants structured the experiment: noun phrase type (revealing or non-revealing) and concept type (noun phrases or nouns). Participants were 32 students at the University of Chicago paid for participating, with 8 assigned randomly to each cell of the noun phrase type X concept type design. None had participated in a previous experiment. Each participant generated properties for 10 target concepts and for 10 filler concepts, received in one of 3 random orders, distributed equally across conditions.

9.1.2. Materials

Table B3 in Appendix B presents the critical materials in Experiment 3. Each of the 20 critical nouns referred to an entity whose internal properties are normally not salient (e.g., *lawn*). The same 10 critical modifiers were used in both the revealing and non-revealing noun phrase conditions. In the revealing noun phrase condition, each noun phrase referred to an object whose interior is salient when the modifier concept applies (e.g., *rolled-up lawn*). In the non-revealing noun phrase condition, the noun phrase referred to an object whose interior is not salient when the modifier concept applies (e.g., *rolled-up snake*). Because the two noun phrase conditions shared modifiers but differed in nouns, the 10 nouns in each of the two noun conditions differed. The same 10 filler nouns and the same 10 filler noun phrases from Experiment 2 were used here in the respective noun and noun phrase conditions.

9.1.3. Procedure

The procedure was the same as in Experiments 1 and 2 except that all participants received neutral instructions.

9.2. Results

The 10 critical protocols of 7 randomly selected participants were coded by two independent raters. Because agreement was 92%, only one judge coded the remainder of the protocols.

Table 7
Proportions of internal and external properties for nouns, noun phrases with revealing modifiers, and noun phrases with non-revealing modifiers in Experiment 3.

Property	Revealing modifiers		Non-revealing modifiers	
	Nouns	Noun Phrases	Nouns	Noun phrases
Internal	.15	.37	.08	.07
External	.40	.21	.50	.38

9.2.1. Overall performance

Participants produced 10.79 properties per concept on average. Nouns (14.06) produced more properties than noun phrases (7.51) ($F'(1, 28) = 6.45$, $MS_e = 14.40$, $p < .01$). Revealing nouns and noun phrases (11.63) did not differ from non-revealing nouns and noun phrases in the number of properties generated (9.94) ($F'(1, 31) = 0.40$, $MS_e = 15.43$). These two factors did not interact ($F'(1, 27) = 0.07$, $MS_e = 14.02$).

9.2.2. Occlusion

The simulation account predicts that occlusion should only affect the production of internal properties for revealing noun phrases, not for non-revealing noun phrases. For revealing noun phrases, internal properties should be relatively inaccessible for nouns when they are occluded but should become more accessible for noun phrases when they are unoccluded. For non-revealing noun phrases, internal properties should not become more accessible for noun phrases. Because internal properties do not become unoccluded in the simulations that represent non-revealing noun phrases, the accessibility of internal properties should not increase. Conversely, if rules are associated with the modifiers in these noun phrases, they should make internal properties more accessible when combined with these nouns, just as they do when combined with any other noun, including those in the revealing noun phrase condition.

As in Experiments 1 and 2, the proportion of internal properties was computed for each protocol that a participant produced for a concept, as was the proportion of external properties. Table 7 presents the average proportions from this analysis. Similar to Experiments 1 and 2, a planned comparison assessed the predictions that the proportion of internal properties should be higher for revealing noun phrases than for nouns, whereas the proportion of external properties should be higher for nouns than for revealing noun phrases. Conversely, it was predicted that the analogous contrast should not be significant for the non-revealing noun phrases. Thus, one contrast was computed for revealing noun phrases relative to nouns, and a second was computed for non-revealing noun phrases relative to nouns. Each contrast—applied individually to a protocol—again used the following weights: internal properties for nouns phrases (+1), internal properties for nouns (-1), external properties for noun phrases (-1), and external properties for nouns (+1).

In support of the simulation account, only the contrast for the revealing noun phrases relative to their nouns was significant across participants and protocols ($t_{DS}(4,138) = 5.25$, $MS_e = .67$, $p < .001$). Conversely, the contrast for non-revealing noun phrases relative to their nouns was not significant ($t_{DS}(4,138) = 1.60$, $MS_e = .67$). As predicted, occlusion modulated the production of internal and external properties in an inversely related manner for revealing noun phrases but not for non-revealing noun phrases. Thus, simulation appears responsible for occlusion effects, not rules.

Specific contrasts within these interactions further supported the predictions of the simulation account. For revealing noun phrases, the proportion of internal properties was higher for noun phrases (.37) than for nouns (.15), ($t_{DS}(2,138) = 3.99$, $MS_e = .67$, $p < .001$). Conversely, for non-revealing noun phrases, the proportion of internal properties did not differ between noun phrases (.07) and nouns (.08), ($t_{DS}(2,138) = .10$, $MS_e = .67$). As the simulation account further predicts for revealing noun phrases, the proportion of external properties was higher for nouns (.40) than for noun phrases (.21) ($t_{DS}(2,138) = 3.39$, $MS_e = .67$, $p < .001$). For the non-revealing noun phrases, the proportion of external properties (.50) was marginally higher for nouns than for noun phrases (.38) ($t_{DS}(2,138) = 2.19$, $MS_e = .67$, $p < .10$).

9.2.3. Situated simulation

Table 8 illustrates that types of properties again differed in frequency ($F'(3,84) = 26.47$, $MS_e = 1.86$, $p < .0001$). Entity properties (.40) were more frequent than situational properties (.35) ($t_{DS}(3,84) = 5.00$, $MS_e = 1.86$, $p < .01$). In addition, situational properties were more frequent than introspective properties (.15), which were more frequent than taxonomic categories (.02) ($t_{DS}(3, 84) = 20.03$, $MS_e = 1.86$, $p < .0001$; $t_{DS}(3,84) = 12.98$, $MS_e = 1.86$, $p < .001$, respectively). Again, it is not surprising the entity properties were produced most often, given that participants were asked explicitly to produce them. Of greater interest is the finding that participants situated their simulations of nouns and noun phrases. Consistent with the hypothesis that object simulations are situated, 50% of the properties described situations and introspections—more than the proportion that described entities (.40)! As in Experiments 1 and 2, these data strongly suggest that participants situated their object conceptualizations with respect to physical settings and mental perspectives.

9.3. Discussion

As in Experiments 1 and 2, the results support the a priori predictions of the simulation account. Again, occlusion had strong effects on property generation. When internal properties were salient for the revealing noun phrases, participants reported them at higher rates than when they were not salient for their respective nouns. As internal properties become more frequent, external properties became less frequent.

Most importantly, the non-revealing noun phrases did not exhibit this pattern of results—internal properties did not become more

Table 8
Proportions and frequencies of general property types in Experiment 3.

Property	Proportion	Frequency			
		Revealing modifiers		Non-revealing modifiers	
		Nouns	Noun phrases	Nouns	Noun phrases
Taxonomic	.02	.65	.09	.34	.13
Entity	.40	6.11	4.38	5.15	1.94
Situation	.35	5.53	2.48	4.59	2.56
Introspection	.15	1.58	1.09	2.31	1.35
Total		14.58	8.68	13.55	6.34

Note: Because miscellaneous properties are not included above, the frequencies of category, entity, situation, and introspection properties do not sum to the total number of properties, nor do their proportions sum to 1.

accessible for these conceptual combinations. This finding indicates that revealing modifiers do not have rules associated with them that generally increase the accessibility of internal properties in the nouns they modify. If they did, the accessibility of internal properties should have increased for both the revealing and non-revealing noun phrases. In contrast, the simulation account explains this finding. When internal properties remained occluded in simulations that represented the non-revealing noun phrases, they did not increase in accessibility.

Traditional theories could attempt to explain these findings by positing an individual rule for every combination of a modifier and a noun. Thus, one rule could state that internal properties increase in importance when *rolled-up* modifies *lawn*, whereas a second rule could state that internal properties do not change in importance when *rolled-up* modifies *snake*. One problem with this account is that, first, it is completely post hoc. Nothing in traditional theories predicts a priori that such rules should exist. A second problem is that the individual rule account does not explain how these rules arise in the first place. How does the cognitive system know to construct one rule that increases the importance of internal properties for *rolled-up lawn* but to construct a different rule that leaves their importance for *rolled-up snake* unchanged? Clearly it is perceptual representations that motivate these rules! Whether internal properties become more or less accessible depends on whether they are salient or occluded in a perceptual representation. Thus, perceptual simulations provide a more parsimonious account of these results than rules, given that simulations can explain these results on their own, whereas rules cannot. A final problem for the individual rule account is its implausibility. An infinite number of amodal rules would be required to represent the infinite number of conceptual combinations that people can construct. Instead, information about occlusion appears to reside implicitly in the simulations that underlie these combinations.

A related argument applies to the results of Experiment 2. How did participants know that internal properties should be more important for the novel noun phrases (e.g., *glass car*) than for their respective nouns (e.g., *car*)? Because participants had presumably not seen referents of these novel noun phrases, they lacked perceptual experiences that could be recoded into rules. Instead, participants probably simulated these concepts perceptually (by combining simulations for the modifiers and head nouns), thereby causing internal properties to become salient. It is possible that individual rules could have been constructed subsequently to recode these simulations, but assuming that participants simply relied on perceptual simulations is more parsimonious.

Finally, as in the previous experiments, the results indicated that participants situated the representations they used to generate properties. Indeed situational and introspective properties constituted 50% of participants' protocols, more than entity properties (40%). Participants spent more time describing background situations than describing target objects.

10. General discussion

We first summarize results across experiments. We then assess implications of these results for the process of property generation and for the process of conceptual combination.

10.1. Summary of results

10.1.1. Occlusion

Across all three experiments, occlusion modulated the proportion of internal and external properties that participants produced. When internal properties were occluded in the noun condition, their accessibility was lower than when they were unoccluded in the noun phrase condition. Conversely, as internal properties became more accessible in the noun phrase condition, external properties became less accessible.

These findings support the a priori prediction of the simulation account that participants construct simulations to represent the meanings of noun phrases. Because these simulations contain visual representations, occlusion affects processing. When internal properties are occluded behind an external surface, they are relatively inaccessible and therefore not reported frequently. When the occluding surface is removed, internal properties become visible and are reported more often. Conversely, as the internal region of an object fills more of a simulation, the object's external surface fills it less, such that external properties are generated less often. Planned comparisons that assessed this inverse relationship between internal and external properties were significant across all three experiments.

10.1.2. Instructional equivalence

As the simulation account further predicts, neutral and imagery participants produced highly similar distributions of properties that differed from those for word participants. After correcting for attenuation due to unreliability, the correlations between neutral and imagery participants were nearly perfect. Furthermore, partial correlations indicated that imagery performance explained nearly all the variance in neutral performance, whereas word performance explained little. Together these results indicate that neutral participants spontaneously constructed simulations—much like those of imagery subjects—to produce properties.

Interestingly, however, neutral participants produced significantly fewer properties than imagery participants. Whereas imagery instructions may have induced conscious imagery that was relatively rich in detail, neutral instructions may have induced imagery that was more unconscious and sketchier. Nevertheless, both groups exhibited occlusion effects and instructional equivalence. Although the amount of perceptual detail in simulations appeared to vary between the neutral and imagery groups, much underlying perceptual structure appeared to remain constant.

10.1.3. Situated simulation

Across experiments, participants produced substantial numbers of properties that described background situations for the target concepts. The total proportion of these properties was 26%, 35%, and 50% across Experiments 1, 2, and 3, respectively. Indeed participants in Experiment 3 produced more situational properties (50%) than entity properties (40%). Situational properties described settings and events, ranging from 19% to 35% across experiments, and described introspections, ranging from 6% to 15%. Again, the simulation account explains these results. Just as attended objects exist against a background situation in perception, conceptualized objects exist against a background situation in conceptualization. In neither are focal objects represented in isolation. Because the conceptual system attempts to simulate perception, it represents simulated objects in simulated situations.

10.1.4. Alternative accounts

Experiment 2 addressed the possibility that simulations are only used to represent familiar conceptual combinations whose referents have been experienced (e.g., *convertible car*) but not for novel conceptual combinations that must be represented with predicate-calculus-like structures (e.g., *glass car*). If this account were correct, then occlusion effects and instructional equivalence should only occur for familiar combinations, not for novel ones. In Experiment 2, however, occlusion effects and instructional equivalence occurred for novel combinations. These results support the simulation account, which proposes that people combine simulations for individual concepts to represent novel combinations (Barsalou, 1999, 2003, 2005a, 2008b).

Experiment 3 addressed the possibility that rules associated with revealing modifiers produce occlusion effects by lowering the accessibility of internal properties in predicate-calculus-like structures. If this account were correct, then occlusion effects should generally occur whenever these modifiers are coupled with a noun in a noun phrase. In Experiment 3, however, these modifiers did not produce occlusion effects for nouns whose internal regions are not revealed (e.g., *rolled-up snake*). These results indicate that revealing modifiers are not associated with general rules that decrease the accessibility of internal properties. Instead, occlusion effects result from internal properties becoming visible in simulations (e.g., *rolled-up lawn*).

10.2. Toward an account of property generation

These experiments suggest that participants construct simulations to support property generation. When participants receive a noun or noun phrase, they construct a simulation to represent it, scan across the simulation, and describe properties perceived in the simulation. Because simulations are scanned and described in this manner, unoccluded properties are described relatively often, whereas occluded properties are described less. Because neutral participants construct simulations, they produce similar properties to those of imagery participants. Because imagery participants were instructed explicitly to construct images, their images are richer, and as a result, have more to describe, thereby producing more properties than neutral participants. Results from Simmons, Hamann, Harenski, Hu, and Barsalou (2008) and from Santos et al. (2008) further support the conclusion that participants use simulations to generate properties.

Barsalou (2003a), (2005a) offers an account of how participants produce properties from simulations. When participants are asked to generate properties for a concept, they retrieve a perceptual memory of an instance, or of a type of instance (prototype), encountered frequently. At this point in processing, the retrieved instance is represented as a relatively holistic unparsed image that is likely to exist on multiple modalities, including modalities in perception, action, and introspection. As the participant begins focusing attention on regions of the multimodal simulation, simulators become active that construe (interpret) these regions. A simulator is analogous to a concept in traditional theories but is implemented using simulation mechanisms (Barsalou, 1999). More specifically, a simulator aggregates memories of modal states across perception, action, and introspection from previous experiences with a concept's instances, where this aggregate information is comparable to the aggregate information in the exemplar and prototype constructs of traditional theories. Once aggregate multi-modal information exists in a simulator, it implements classic functions of concepts.

Most importantly for our purposes here, simulators implement type-token binding. Imagine that a participant has been asked to produce properties for *car* and has simulated a particular car from memory. As the participant focuses attention on a region of the simulation such as a wheel, the simulator for *wheel* becomes active and then bound to the wheel in the simulation (and possibly to other wheels). Because the *wheel* simulator contains visual information captured from previous wheels, it bears sufficient similarity to the wheel in the simulation to become active. Once the *wheel* simulator has been bound to the wheel in the simulation, it implements the functionality of a type-token proposition: Aggregated information about the simulator's instances has become bound to a particular instance, construing it as a wheel.

Such construals can be true or they can be false, thereby implementing classic propositions (Barsalou, 1999). Furthermore, information in the simulator that is not in the instance can be extended to the instance as an inference. For example, if the instance of a wheel in the simulation is static, the simulator can generate an inference that the wheel could roll—implemented as a simulation of the wheel rolling—based on aggregated information acquired from seeing rolling wheels. Many further inferences from multiple modalities could similarly be generated through simulation, such as how a tire sounds, feels, and smells (e.g., Gil-da-Costa et al., 2004; Simmons, Martin, & Barsalou, 2005).

Once a simulator becomes bound to a region of a simulation to implement a type-token proposition, the word for the simulator is generated as a property of the concept. For example, once the *wheel* simulator is bound to a simulated wheel, the associated word "wheel" is generated. Because many possible simulators—and combinations of embedded simulators—could in principle be bound to regions of a car, tremendous diversity is possible in describing it. Not only can simulators for isolated properties become bound, so can simulators for relations and configurations of simulators that represent propositions (Barsalou, 1999, 2003a, 2005a). Additionally, when simulators for settings, events, and introspections become active to describe the background situation, situational construals occur.

The diversity possible in describing a simulation implements a type of creativity often observed for property generation that Barsalou (1993) referred to as "linguistic vagary." Such diversity is also a likely source of the substantial variability that occurs between different participants generating properties for the same concept (only 44% overlap on the average), and that similarly occurs for the same participant generating properties on different occasions (only 67% overlap on the average); see Barsalou (1989), (1993) for further discussion. Because participants simulate different instances of concepts and use different sets of simulators to construe them, variability results in the descriptions produced.

Finally, recent evidence indicates that word association contributes to property generation. Santos et al. (2008) found that classic word associations often constitute the properties generated initially for a concept. When producing properties for *taxi*, the word "cab" is often generated initially, because "cab" is a common word associate that originates from retrieving the compound phrase "taxi cab." In a companion neuroimaging experiment, Simmons et al. (2008) similarly found that word association areas in the brain were active early during property generation. Together these findings suggest that a combination of word association and situated simulation is responsible for generating the properties of concepts (Barsalou et al., 2008).

10.3. Toward an account of conceptual combination

The literature on conceptual combination contains many empirical findings and many accounts of these findings. Notably, these accounts almost always adopt predicate-calculus-like structures to represent concepts. It might therefore seem that the simulation results reported here, and the account of them just presented, would be at odds with other existing accounts of conceptual combination. To the contrary, we believe that, with one major exception, our findings and account can be integrated with other findings and accounts.

The major exception concerns representation. Whereas traditional theories of conceptual combination assume that predicate-calculus-like structures represent concepts, we assume that simulators representation concepts instead. As just described, however, the simulation account implements the functionality of predicate-calculus-like representations. Thus, the major difference between accounts is how this functionality is implemented—not its importance. Occlusion effects and instructional equivalence suggest that simulations implement this functionality. These effects, however, do not imply that this functionality does not exist. To the contrary, we strongly believe that it does for all the reasons cited in Barsalou (1999), (2003a), (2005a), (2008b).

Because the simulation account incorporates the functionality of traditional theories, it is compatible with the effects reported in the conceptual combination literature, and with the functionality (but not the representations) of the theories developed to explain them. Consider some examples. Smith et al. (1988) argued that the modifier in a noun phrase instantiates an argument in a predicate for the head noun, overriding the default value. In *blue house*, for example, *blue* instantiates the *color* argument in the *house* predicate, overriding the default value of *white*. The simulation account implements this by assuming that a house is simulated in a default color (white) and then transformed visually to blue, as specified by modifier. Functionally, a *color* argument is being bound to different values for *white* and *blue*, but the implantation occurs at the level of simulators operating on a simulation, not through changes in a predicate-calculus-like structure.

Medin and Shoben (1988) showed that instantiating an argument in a predicate typically propagates correlated values to other arguments. For example, *large bird* does not simply change the value of *size* for *bird* but also changes the value of *beak* from *straight* to *curved* and the value of *vocalization* from *sings* to *screeches*. The simulation account explains this in two ways. First, the holistic simulations retrieved initially to represent a concept implicitly contain correlated properties (e.g., memories of large birds tend to contain curved beaks and screeching). Second, simulators for these values become active to interpret the holistic simulation (as described above). Because these simulators tend to be correlated in memory through past experience, when one simulator construes a simulation (e.g., *large*), correlated simulators are retrieved as inferences (e.g., *curved beak* and *screeches*).

Many researchers have argued that various types of relations between properties are central to conceptual combination (e.g., Gagne & Shoben, 1997; Wisniewski, 1997). In noun-noun combinations, for example, the modifying noun can be related to the head noun by a location relation (e.g., *mountain flower*), an instrument relation (e.g., *phone message*), an object relation (e.g., *bicycle repair*), or by a variety of other relations. In the simulation account, these relations can be implemented with relation simulators (Barsalou, 1999, 2003a, 2005a). For example, the *on* simulator is acquired by processing an object in a particular form of contact with another object across many instances (e.g., a car on a driveway, a fly on a wall, and a flower on a mountain). In a simulation of a flower on a mountain, for example, the *on* simulator becomes active and binds its region for a focal object to *flower* and its region for a background object to *mountain*. Furthermore, the fact that many of the important relations in conceptual combination are thematic is consistent with the finding that simulations tend to be situated (Experiments 1, 2, and 3 here; Barsalou, 2003b; Barsalou, 2005b; Santos et al., 2008; Simmons et al., 2008; Yeh & Barsalou, 2006). Because simulations are typically situated, the thematic relations that exist between situation components are implicitly available in them, and can be represented explicitly through the application of relation simulators. We assume that much of the background knowledge important for conceptual combination arises through these same mechanisms (e.g., Murphy, 1988).

Many researchers have shown that the combination of two concepts is not a simple union of their properties (for a review, see Hampton, 1991). Instead interactions typically occur that cause some properties in the union to become irrelevant and other properties not in the union to become relevant. The simulation account offers two explanations of these non-linearities. First, the holistic simulation retrieved initially for a conceptual combination may differ from the holistic simulations retrieved for the individual concepts. If so, then the simulation for the combination offers different opportunities for the simulator-binding process described above that construes simulations. As a result, different properties become salient for the combination than for the individual concepts. Second, once these different simulators become bound to the simulation, they may activate associated simulators as inferences, not relevant for the individual concepts.

Consider the classic example of *pet fish* (Hampton, 2007). *Pet* tends to activate simulations of dogs and cats, and *fish* tends to activate simulations of ocean and lake fish. In contrast, *pet fish* tends to activate simulations of *goldfish* and *guppies.* Because *pet fish* tend to be small and docile, simulations of them tend to be construed differently than simulations of ocean and river fish, which tend to be large and wild. Furthermore, because *pet fish* swim, they tend to be described differently than *pet*, which is typically assumed to walk. Additionally, the correlated networks of simulators associated with *pet, fish,* and *pet fish* differ, such that once simulators in one of these networks become active, they tend to activate different associated simulators as inferences, thereby producing further differences in how the individual and combined concepts are construed.

Finally, Costello and Keane (2000) proposed that pragmatic constraints such as plausibility, diagnosticity, and informativeness operate when people construct one of the many interpretations possible for a noun phrase. Notably, they argue that their account is functional and does not imply a particular form of representation (e.g., predicate calculus vs. simulation). Furthermore, they argue that background knowledge is central for implementing these constraints. As just described, we assume that background knowledge is carried implicitly in simulations of situations and explicitly in the correlated networks of simulators used to construe them. Thus, the simulation account appears capable of providing a framework within which these Costello and Keane's constraints can be implemented.

Clearly these sketches of how the simulation account explains the conceptual combination literature are little more than promissory notes. Our point, however, is that the simulation account is not necessarily at odds with findings in this literature, nor with the functionality of the accounts used to explain them. We further suspect that the simulation account has the ability to strengthen and support previous accounts. Because the simulation account assumes that tremendous amounts of information from perception, action, and introspection are stored in the conceptual system, background knowledge and relations central to conceptual combination reside in this information implicitly. Indeed a problem with existing theories is that the predicate-calculus structures for a concept often seem unmotivated and contrived, with it not being clear why these particular structures and not others exist (Barsalou, 1993, 2003a, 2005a). Because the simulation account grounds these structures in perceptual, motor, and introspective experience, and constructs them from this experience, it offers an approach for explaining them in a principled manner.

Acknowledgements

We are grateful to Sarah Bunton for laboratory assistance, to Karen Solomon, Wenchi Yeh, and Barbara Luka for helpful discussion, and to Edward Wisniewski and Karen Solomon for helpful comments on earlier drafts. This work was supported by National Science Foundation Grants SBR-9421326, SBR-9796200, SBR-9905024, and BCS-0212134 to Lawrence W. Barsalou. Address correspondence to Ling-ling Wu, Department of Information Management, National Taiwan University, Taipei, Taiwan (llwu@management.ntu.edu.tw), or to Lawrence W. Barsalou, Department of

Psychology, Emory University, Atlanta, GA 30322 (barsalou@emory.edu), <http://www.psychology.emory.edu/cognition/barsalou/index.html>).

Appendix A. General and specific codes for generated properties

In the coding categories below, lowercase text in italics refers to concepts, whereas uppercase text refers to a coded property of the concept. Note that more recent versions of this coding scheme have been developed.[4]

Taxonomic categories (C). A category in the taxonomy to which a concept belongs.

Synonym (C_S). A synonym of a concept (e.g., *car*–AUTOMOBILE; *cat*–FELINE).
Ontological category (C_O). A category for a basic kind of thing in existence, including *thing, substance, object, human, animal, plant, location, time, activity, event, action, state, thought, and emotion* (e.g., *cat*–ANIMAL; *computer*–OBJECT).
Superordinate (C_H). A category one level above a concept in a taxonomy (e.g., *car*–VEHICLE; *apple*–FRUIT).
Coordinate (C_C). Another category in the superordinate category to which a concept belongs (e.g., *apple*–ORANGE; *oak*–ELM).
Subordinate (C_L). A category one level below the target concept in a taxonomy (e.g., *chair*–ROCKING CHAIR; *frog*–TREE FROG).
Individual (C_I). A specific instance of a concept (e.g., *car*–MY CAR; *house*–MY PARENTS' HOUSE).

Entity Properties (E). Properties of a concrete entity, either animate or inanimate. Besides being a single self-contained object, an entity can be a coherent collection of objects, or an institution, if it consists of at least some concrete entities (e.g., *forest, government*, and *society*).

Larger whole (E_W). A whole to which an entity belongs (e.g., *window*–HOUSE; *apple*–TREE).
Spatial relation (E_S). A spatial relation between two or more properties within an entity, or between an entity and one of its properties (e.g., *car*–window ABOVE door; *watermelon*–green OUTSIDE).
External surface property (E_{SE}). An external property of an entity that is not a component, and that is perceived on or beyond the entity's surface, including *shape, color, pattern, texture, size, touch, smell, and taste* (e.g., *watermelon*–OVAL; *apple*–RED; *car*–STINKS).
Internal surface property (E_{SI}). An internal property of an entity that is not a component, that is not normally perceived on the entity's exterior surface, and that is only perceived when the entity's interior surface is exposed, including *color, pattern, texture, size, touch, smell, and taste* (e.g., *apple*–WHITE, *watermelon*–JUICY).
External component (E_{CE}). A three-dimensional component of an entity that, at least to some extent, normally resides on its surface (e.g., *car*–HEADLIGHT; *tree*–LEAVES).
Internal component (E_{CI}). A three-dimensional component of an entity that normally resides completely inside the closed surface of the entity (e.g., *apple*–SEEDS; *jacket*–LINING).
Systemic property (E_{SYS}). A global systemic property of an entity or its parts, including *states, conditions, abilities, and traits* (e.g., *cat*–ALIVE; *dolphin*–INTELLIGENT; *car*–FAST).
Entity behavior (E_B). An intrinsic action that is characteristic of an entity's behavior, and that is not an entity's normal function for an external agent, which is coded as S_F (e.g., *dog*–BARKS; *children*–PLAY).
Associated abstract entity (E_{AE}). An abstract entity associated with the target entity and external to it (e.g., *computer*-SOCIETY; *transplanted Californian*-RELIGIOUS AFFILIATION).
Quantity (E_Q). A numerosity, frequency, or intensity of an entity or its properties (e.g., *jacket*–an ARTICLE of clothing; *cat*–FOUR legs; *tree*–LOTS of leaves; *apple*–COMMON fruit; *watermelon*–USUALLY green; *apple*–VERY red).

Situation Properties (S). A property of a situation, where a situation typically includes one or more participants, at some place and time, engaging in an event, with one or more entities (e.g., *picnic, conversation, vacation, and meal*).

Participant (S_P). A person in a situation who typically uses an entity or performs an action on it and/or interacts with other participants (e.g., *toy*–CHILDREN; *car*-PASSENGER; *furniture*–PERSON).
Location (S_L). A place where an entity can be found, or where people engage in an event or activity (e.g., *car*–IN THE GARAGE; *buy*–IN A STORE).
Spatial relation (S_S). A spatial relation between two or more things in a situation (e.g., *watermelon*–the ants crawled ACROSS the picnic table; *car*–drives ON the highway; *vacation*–we slept BY the fire).
Time (S_T). A time period associated with a situation or with one of its properties (e.g., *picnic*–FOURTH OF JULY; *sled*–DURING THE WINTER).
Action (S_A). An action that a participant performs in a situation (e.g., *shirt*–WORN; *apple*-EATEN).
Associated entity (S_E). An entity in a situation that contains the focal concept (e.g., *watermelon*–TABLE; *cat*–LITTER).
Function (S_F). A typical goal or role that an entity serves for an agent (e.g., *car*–TRANSPORTION; *clothing*–PROTECTION).
Quantity (S_Q). A numerosity, frequency, or intensity of a situation or any of its properties except of an entity, whose quantitative aspects are coded with E_Q (e.g., *vacation*–lasted for EIGHT days; *car*–a LONG drive).
Manner (S_M). The manner in which an action or behavior is performed (e.g., *watermelon*–SLOPPY eating; *car*–FASTER than walking).
Event (S_{EV}). An event or activity in a situation (e.g., *watermelon*–PICNIC, *car*–TRIP).
State of the world (S_{SW}). A state of a situation or any of its components except entities, whose states are coded with E_{SYS} (e.g., *mountains*–DAMP; *highway*–CONGESTED).
Origin (S_{OR}). How or where an entity originated (e.g., *car*–FACTORY; *watermelon*–GROUND).

Introspective properties (I). A property of a participant's mental state as he or she views a situation, or a property of a participant's mental state in a situation.

Affect/emotion (I_A). An affective or emotional state toward the situation or one of its components by either the participant or a participant (e.g., *magic*–a sense of EXCITEMENT; *vacation*–I was HAPPY; *smashed car*–ANGER).
Evaluation (I_E). A positive or negative evaluation of a situation or one of its components by either the participant or a participant (e.g., *apples*–I LIKE them; *vacation*–I wrote a STUPID paper).

[4] For more recent versions, contact the second author or the authors of McRae et al. (2005) and Brainerd et al. (2008).

Representational state (I_R). A representational state in the mind of a situational participant, including *beliefs, goals, and ideas* (e.g., *smashed car*–believed it was not working; *cut tree*–wanted to cut it down).

Cognitive operation (I_O). An operation on a cognitive state, including *retrieval, comparison, and learning* (e.g., *watermelon*–I REMEMBER a picnic; *rolled up lawn*–LOOKS LIKE a burrito; *car*–I LEARNED how to drive).

Contingency (I_C). Contingency between two or more aspects of a situation, including *if, enable, cause, because; depends, requires, correlated, uncorrelated, and negatively correlated* (e.g., *car*–REQUIRES gas; *tree*–has leaves DEPENDING ON the type of tree; *vacation*–FREE FROM work; *magic*–I was excited BECAUSE I got to see the magician perform).

Quantity (I_Q). Numerosity, frequency, or intensity of an introspection or one of its properties (e.g., *truth*–a SET of beliefs; *buy*–I was VERY angry at the saleswoman; *magic*–I was QUITE baffled by the magician).

Negation (I_N). An explicit mention of the absence of something, with the absence requiring a mental state that represents the opposite (e.g., *car*–NO air conditioning, *apple*–NOT an orange).

Miscellaneous (M). Information in a protocol not of theoretical interest.

Repetition (M_R). Repetition of an item already coded (e.g., *car*–has wheels, WHEELS; *apple*–round, ALMOST ROUND)

Meta-comment (M_M). A meta-comment about the task (e.g., *house*–THEY CAN TAKE SO MANY FORMS; transplanted Californian–IT IS HARD TO IMAGINE THIS).

Appendix B

Table B1
Critical materials in Experiment 1.

Noun	Noun phrase
Lawn	Rolled-up lawn
Watermelon	Half watermelon
Apple	Sliced apple
Banana	Peeled banana
Tomato	Smashed tomato
Tree	Cut tree
Face	Smiling face
Sheep	Shaved sheep
Jacket	Inside-out jacket
Mattress	Torn mattress

Table B3
Critical materials in Experiment 3.

Revealing		Non-revealing	
Nouns	Noun phrases	Nouns	Noun phrases
Lawn	Rolled up lawn	Snake	Rolled up snake
Watermelon	Half watermelon	Smile	Half smile
Tomato	Smashed tomato	Car	Smashed car
Arm	Cut arm	Lawn	Cut lawn
Computer	Open computer	Hands	Open hands
Bus	Empty bus	Sofa	Empty sofa
Car	Comfortable car	Watch	Comfortable watch
Dresser	Disorganized dresser	Plumber	Disorganized plumber
Chapel	Candle-lit chapel	Priest	Candle-lit priest
Bush	Transplanted bush	Californian	Transplanted Californian

Table B2
Critical materials in Experiment 2. rated for phrase and referent familiarity.

Novel combinations			Familiar combinations		
Combination	Phrase	Referent	Combination	Phrase	Referent
Rolled-up lawn	1.31	1.81	Dug-up lawn	2.81	2.75
Chipped coconut	1.18	1.88	Split coconut	2.75	2.50
Gashed watermelon	1.38	1.94	Half watermelon	3.06	3.44
Mangled cow	1.12	1.19	Butchered cow	2.69	2.81
Butchered frog	1.19	1.38	Dissected frog	3.06	2.25
Laughing cat	1.25	1.43	Yawning cat	1.94	2.81
Glass car	1.06	1.19	Convertible car	3.63	3.56
Transparent computer	1.06	1.25	Opened computer	1.5	2.31
Demolished clock	1.50	1.56	Glass clock	1.88	2.69
Inside-out purse	1.38	1.81	Open purse	3.56	4.06
Average	1.24	1.54	Average	2.69	2.92

References

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review, 85*, 249–277.

Barsalou, L. W. (1989). Intra-concept similarity and its implications for inter-concept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. New York: Cambridge University Press. pp. 76–121.

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

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of memories*. London: Erlbaum. pp. 29–101.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–660. Available in pdf format at <http://userwww.service.emory.edu/barsalou/Publications/online.html>.

Barsalou, L. W. (2002). Being there conceptually: Simulating categories in preparation for situated action. In N. L. Stein, P. J. Bauer, & M. Rabinowitz (Eds.), *Representation, memory, and development: Essays in honor of Jean Mandler* (pp. 1–19). Mahwah, NJ: Erlbaum.

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. (Reprinted in Moss, H., Hampton, J. *Conceptual representation* (pp. 513–566). East Sussex, UK: Psychology Press.).

Barsalou, L. W. (2005a). Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe, D. Rakison (Eds.), *Building object categories. Carnegie symposium series* (pp. 389–431). Mahwah, NJ: Erlbaum.

Barsalou, L. W. (2005b). Continuity of the conceptual system across species. *Trends in Cognitive Sciences, 9*, 309–311.

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 (Eds.), *Embodied grounding: Social, cognitive, affective, and neuroscientific approaches* (pp. 9–42). New York: Cambridge University Press.

Barsalou, L. W. (2008c). Situating concepts. In P. Robbins & M. Aydede (Eds.), *Cambridge handbook of situated cognition* (pp. 236–263). 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., Niedenthal, P. M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The psychology of learning and motivation*, (Vol. 43, pp. 43–92). 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 (Eds.), *Symbols, embodiment, and meaning* (pp. 245–283). Oxford: Oxford University Press.

Brainerd, C. J., Yang, Y., Reyna, V. F., Howe, M. L., & Mills, B. A. (2008). Semantic processing in "associative" false memory. *Psychonomic Bulletin and Review, 15*, 1035–1053.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior, 12*, 335–359.

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

Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science, 244*, 299–349.

Cree, G. S., & McRae, K. (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.

Finke, R. A. (1989). *Principles of mental imagery*. Cambridge, MA: MIT Press.

Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.

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

Gagne, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 71–87.

Gil-da-Costa, R., Braun, A., Lopes, M., Hauser, M. D., Carson, R. E., Herscovitch, P., et al. (2004). Toward an evolutionary perspective on conceptual representation: Species-specific calls activate visual and affective processing systems. *Proceedings of the National Academy of Sciences, 101*, 17516–17521.

Glaser, W. R. (1992). Picture naming. *Cognition, 42*, 61–106.

Guilford, J. P., & Fruchter, B. (1973). *Fundamental statistics in psychology and education* (5th ed.). New York: McGraw-Hill.

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

Hampton, J. A. (1991). The combination of prototype concepts. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings*. Hillsdale, NJ: Lawrence Erlbaum Associates. pp. 91, 116.

Hampton, J. A. (2007). Typicality, graded membership, and vagueness. *Cognitive Science, 31*, 355–384.

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

Kirk, R. E. (1982). *Experimental design: Procedures for the behavioral sciences* (2nd ed.). Pacific Grove, CA: Brooks/Cole.

Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods, 37*, 547–559.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20*, 158–190.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science, 12*, 529–562.

Pollack, J. (1990). Recursive distributed representations. *Artificial Intelligence, 46*, 77–105.

Pulvermüller, F., Shtyrov, Y., & Ilmoniemi, R. (2005). Brain signatures of meaning access in action word recognition. *Journal of Cognitive Neuroscience, 17*, 884–892.

Pylyshyn, Z. W. (1984). *Computation and cognition*. Cambridge, MA: MIT Press.

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

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.

Santos, A., Chaigneau, S. E., Simmons, W. K., Barsalou, L. W. (2008). *Word association and situated simulation in conceptual processing: Behavioral confirmation of neuroimaging findings*. Under review.

Shepard, R. N., & Cooper, L. A. (1982). *Mental images and their transformations*. New York: Cambridge University Press.

Simmons, W. K., Hamann, S. B., Harenski, C. N., Hu, X. P., Barsalou, L. W. (2008). fMRI evidence for word association and situated simulation in conceptual processing. *Journal of Physiology – Paris, 102*, 106–119.

Simmons, W. K., Martin, A., & Barsalou, L. W. (2005). Pictures of appetizing foods activate gustatory cortices for taste and reward. *Cerebral Cortex, 15*, 1602–1608.

Smith, E. E. (1978). Theories of semantic memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 6, pp. 1–56). Hillsdale, NJ: Lawrence Erlbaum Associates.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science, 12*, 485–527.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence, 46*, 159–216.

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

Winer, B. J. (1971). *Statistical principles in experimental design*. New York: McGraw-Hill.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin and Review, 4*, 167–183.

Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *American Journal of Psychology, 119*, 349–384.