

Property generation reflects word association and situated simulation

AVA SANTOS⁺, SERGIO E. CHAIGNEAU[~], W. KYLE SIMMONS⁼ AND
LAWRENCE W. BARSALOU^{∞*}

⁺Fort Lewis College

[~]Universidad Adolfo Ibanez

⁼Laureate Institute for Brain Research

[∞]Emory University

Abstract

The property generation task (i.e. “feature listing”) is often assumed to measure concepts. Typically, researchers assume implicitly that the underlying representation of a concept consists of amodal propositions, and that verbal responses during property generation reveal their conceptual content. The experiments reported here suggest instead that verbal responses during property generation reflect two alternative sources of information: the linguistic form system and the situated simulation system. In two experiments, properties bearing a linguistic relation to the word for a concept were produced earlier than properties not bearing a linguistic relation, suggesting the early properties tend to originate in a word association process. Conversely, properties produced later tended to describe objects and situations, suggesting that late properties tend to originate from describing situated simulations. A companion neuroimaging experiment reported elsewhere confirms that early properties originate in language areas, whereas later properties originate in situated simulation areas. Together, these results, along with other results in the literature, indicate that property generation is a relatively complex process, drawing on at least two systems somewhat asynchronously.

Keywords

concepts, categories, lexical semantics, features, feature norms, property generation, mental simulation, situated cognition, word association, linguistic forms

* Correspondence address: Ava Santos, Department of Psychology, 1000 Rim Dr., Fort Lewis College, Durango, CO 81301, USA, or Lawrence W. Barsalou, Department of Psychology, Emory University, Atlanta, GA 30322, USA. E-mail: Santos_a@fortlewis.edu or barsalou@emory.edu. This work was supported by National Science Foundation Grants SBR-9421326, SBR-9796200, and BCS-0212134 to Lawrence W. Barsalou.

1. Introduction

For decades, the property generation task—also known as “feature listing”—has been an important tool for measuring conceptual representations (e.g. Cree and McRae 2003; Hampton 1979; McRae et al. 2005; Rosch and Mervis 1975; Rosch et al. 1976).¹ In this task, researchers present participants with the word for a concept and ask them to verbally generate the concept’s properties. Presented with the word “bird,” for example, participants might produce words for the underlying conceptual properties of *feathers*, *wings*, *fly*, *nest*, *tree* and so forth.² These properties are then interpreted as constituting (or at least contributing to) the concept that underlies the word’s meaning. Researchers across diverse areas, including cognitive psychology, social psychology, consumer psychology, clinical psychology, neuropsychology, and cognitive linguistics, use property generation as an important tool for measuring conceptual content.

Because the property generation task is such a central tool for establishing conceptual content, it is important to understand the cognitive and neural mechanisms underlying it. Surprisingly little research, however, has addressed these mechanisms directly. Instead, researchers tend to assume implicitly that property generation measures a unitary form of conceptual representation that takes the form of amodal propositions. Whereas Fodor (1975) and Pylyshyn (1984) make this assumption explicitly, many other theorists make it implicitly (e.g. Collins and Loftus 1975; Murphy 2002; Smith 1978; Smith and Medin 1981).

In this article, we propose, first, that property generation reflects not one but two cognitive processes, and second, that these processes are word association and situated simulation. Should this proposal be correct, it would by no means invalidate previous empirical efforts to collect property norms for concepts—it would simply offer further insight into nature of these norms. Property generation produces useful information about concepts that predicts conceptual processing in diverse tasks. By understanding how people produce properties, we are likely to better understand conceptual processes in general, and to become even more skilled at predicting conceptual processing and associated behaviors.

-
1. The term “feature listing” is problematic on two counts. First, much of the information produced does not constitute “features,” such as taxonomic categories, events, relations, and so forth. Second, the production task often consists of more than simply “listing” properties, with participants describing complex conceptual structures, using complex syntactic expressions, organizing properties in various manners, and so forth.
 2. Throughout this article we use quotes to indicate the word or phrase for a concept (e.g. “bird”), and italics to indicate concepts (e.g. *feathers*, *wings*).

1.1. *Previous support*

Much previous research and theory predicts that property generation reflects word association and situated simulation. Since the 1960s, Dual Code Theory has generated an impressive body of empirical support for the view that a language-based system and an image-based system underlie cognition (for reviews, see Paivio 1971, 1986). Evidence for these two systems has accumulated in episodic memory, semantic memory, language, thought, developmental psychology, and individual differences. This extensive body of evidence anticipates the proposal that property generation reflects a linguistic process such as word association and an image process such as situated simulation.

Glaser's (1992) extension of Dual Code Theory anticipates our specific proposal even more closely, proposing that the linguistic system can perform relatively superficial processing of linguistic forms independently of the conceptual system, which is adept at processing images. To support his account, Glaser cites findings from many literatures, including the differential ability of pictures vs. words to access the conceptual system during verification tasks, and the differential ability of pictures vs. words to produce priming and interference.

Chaffin (1997) reports experiments consistent with our proposal and results to follow, but interprets them somewhat differently. In these experiments, participants produced free associations to an unfamiliar vs. familiar word from the same taxonomic category (e.g. "sarsaparilla" vs. "soda") or from the same synonym set ("abscond" vs. "escape"). For unfamiliar words, participants tended to produce properties from definitions, whereas for familiar words, participants tended to produce properties from events. Rather than reflecting a single underlying process, the properties produced reflected two processes that we will propose are word association and the situated simulation.

Our proposal follows most closely from previous research on the property verification task (Kan et al. 2003; Solomon and Barsalou 2004). Whereas property generation can be viewed as a production-oriented recall task, property verification can be viewed as comprehension-oriented recognition task. On a property verification trial, participants receive the name of a concept and the name of a property, and then assess whether the property is true of the concept (e.g. *horse-mane* vs. *horse-wall*). Like property generation, property verification has generally been assumed to operate on conceptual information in a unitary store of amodal representations (e.g. Kosslyn 1976; Smith 1978). In contrast, Solomon and Barsalou (2004) found that participants could either use word association or simulation to verify properties. Kan et al. (2003) corroborated these behavioral results in a neuroimaging experiment using the same materials and design. When task conditions blocked the use of word association, a visual area associated with imagery for concrete words was active during property verification. Conversely, when task conditions allowed word

associations, this visual area was not active. Rather than using a single process to verify properties, participants utilized two different processes as task conditions allowed.

Most recently, Andrews et al. (2009) found that conceptual content is modeled better as two processes—one distributional and one experiential—than as a single unitary process, where their distributional and experiential processes correspond to word association and situated simulation in our framework. Furthermore, Andrews et al. found that these two processes do not operate independently but instead appear closely coupled, consistent with our proposal that word association and situated simulation operate in parallel and have corresponding content (Barsalou et al. 2008). Louwerse (2008) presents related arguments and findings.

Finally, we replicated Experiment 2 from this article in a companion neuroimaging experiment (Simmons et al. 2008). As described later, this companion experiment demonstrated that the brain's word production system is active early during property generation, whereas the brain's situated simulation system is active later. Rather than a single neural system being active during property generation, two different systems are active asynchronously.

1.2. *The LASS theory of conceptual processing*

Barsalou et al.'s (2008) LASS theory—Linguistic And Situated Simulation—motivated both the property generation experiments reported here and Simmons et al.'s (2008) companion neuroimaging experiment. LASS theory also motivated our previous work on property verification (Kan et al. 2003; Solomon and Barsalou 2004), and is currently motivating our research on abstract concepts (Barsalou and Wiemer-Hastings 2005; Wilson-Mendenhall, Barrett et al. 2011; Wilson-Mendenhall, Simmons et al. under review). LASS theory grew out of Paivio's (1971, 1986) Dual Code Theory and Glaser's (1992) extension of it. While these three theories are similar in assuming that multiple systems underlie conceptual processing, their specific assumptions vary (see Barsalou et al. 2008). The following sub-sections summarize LASS theory (see Barsalou et al. 2008 for further detail).

1.2.1. *Linguistic processing.*

On perceiving a word, the linguistic system becomes engaged immediately to categorize the linguistic form (which could be auditory, visual, tactile, etc.). Following Figure 1, we assume that the both the linguistic and simulation systems become active immediately, but that activation in the linguistic system peaks first. Because representations of linguistic forms are more similar to presented words than are simulations of their referents, representations of linguistic forms peak earlier (e.g. following the encoding specificity principle of Tulving and Thomson 1973).




Figure 1. *Hypothesized contributions from the linguistic system (L) and the situated simulation system (SS) during conceptual processing according to the LASS theory. When a word is presented as a cue, word form contributions from the linguistic system precede situated simulations from the simulation system. The height, width, shape, and offset of the two distributions are assumed to vary. In response to different words and task conditions, all these parameters are likely to change (e.g. SS activity could be more intense than L activity). The two distributions in this figure illustrate one of many different forms that activations of the L and SS systems could take, while conforming to the principle that, when a word cues a concept, L activity generally precedes SS activity.*

As a word is being recognized, associated linguistic forms are produced as inferences and as pointers to related conceptual information (where by “linguistic forms” we mean the surface forms of words and not their meanings). In the experiments here, the generation of linguistic forms is realized as the simple process of word association, where a cue word elicits associated words (e.g. “car” elicits “automobile” and “vehicle”). We hasten to add that word association is the simplest possible form of the linguistic processing that could occur during conceptual processing. More complex processing occurs as phrases and syntactic structures become active (not addressed here). Furthermore, we assume that a wide variety of relations could exist between a cue word and the associated word generated by it. Of primary relevance here, however, is the proposal that a cue word elicits other linguistic forms during its recognition.

Once associated linguistic forms become active, they support various processing strategies. For example, associations between words can be sufficient to produce correct responses on conceptual tasks, with the activation of conceptual information being unnecessary (e.g. Barsalou et al. 2008; Glaser 1992; Kan et al. 2003; Solomon and Barsalou 2004). Consistent with linguistic context theory (e.g. Burgess and Lund 1997; Landauer and Dumais 1997), active word associates can support a wide variety of current processing demands. Such linguistic strategies are highly consistent with Andrew et al.’s (2009) results that distributional information about language is central to conceptual processing.

These linguistic strategies may be relatively superficial (Glaser 1992). Instead of providing deep conceptual information, these strategies may produce relatively shallow representations that make correct performance possible with

minimal processing. When linguistic forms are sufficient for adequate performance, retrieval of conceptual information is not necessary. Given the obvious heuristic value of such strategies, they can be very useful (cf. Gigerenzer 2000). Mistakenly attributing conceptual depth to these heuristics, however, mischaracterizes them and obscures other important mechanisms that provide deeper conceptual representations.

Much research on lexical processing further supports this proposal. In the lexical decision task, activation of a word's meaning is shallow when the word is read in the context of non-words that violate phonological and orthographic rules; conversely, when non-words satisfy phonological and orthographic rules, words access meaning more deeply (e.g. James 1975; Joordens and Becker 1997; Shulman and Davidson 1977; Stone and Van Orden 1993; Yap et al. 2006). Similarly, the depth-of-processing literature shows that phonemic orienting tasks produce shallower activations of meaning than semantic orienting tasks (e.g. Craik 2002; Craik and Lockhart 1972; Craik and Tulving 1975; Lockhart 2002; Morris et al. 1977). As these findings and others illustrate, linguistic forms can be processed superficially.

1.2.2. *Situated simulation.* As the linguistic system begins to recognize a presented word, the word's representation begins to activate correlated simulations in the brain's modal systems. As Figure 1 illustrates, we assume that the simulation system becomes active very quickly once the presented word form is recognized, but that the activation of a simulation proceeds more slowly than the activation of associated linguistic forms. By "simulation" we mean that the brain simulates the perceptual, motor, and introspective states active during interactions with the word's referents (e.g. Barsalou 1999, 2008a; Damasio 1989; Glenberg 1997; Martin 2001, 2007; Thompson-Schill 2003). Recognizing the word "cat," for example, reenacts neural states that represent how cats look, sound, and feel, how one interacts with cats, and how one feels affectively.

We further assume that simulations are usually situated, thereby preparing agents for situated action (e.g. Barsalou 2003b, 2005b, 2008c; Barsalou et al. 2003; Yeh and Barsalou 2006). Instead of representing the meaning of a word generically, simulations represent them in situations. For example, simulations associated with "cat" do not typically represent cats generically, but instead represent specific cats in particular situations, where a situation contains a setting, agents, objects, actions, events, and mental states.

Finally, we assume that simulations typically represent deep conceptual information, unlike the activation of linguistic forms. We similarly assume that basic symbolic processes including predication, conceptual combination, and recursion result from operations on simulations. Barsalou (1999, 2003a, 2005a, 2008b) describes how simulation mechanisms implement symbolic opera-

tions. We suspect that linguistic mechanisms are not capable of implementing symbolic operations on their own, given that they simply manipulate linguistic forms, not their meanings. Nevertheless, linguistic mechanisms play central roles in controlling simulation during symbolic operations (Barsalou 2008b, Barsalou et al. 2008).

1.2.3. *Mixtures of language and situated simulation.* Different mixtures of the language and simulation systems underlie a wide variety of tasks (Paivio 1971, 1986). When superficial linguistic processing is sufficient to support task performance, processing relies on the linguistic system and little on simulation (Glaser 1992; Kan et al. 2003; Solomon and Barsalou 2004). Conversely, when linguistic processing cannot produce adequate performance, the simulation system supports the required conceptual processing. Depending on task conditions, conceptual processing may primarily depend on either linguistic processing or on simulation. Under many conditions, conceptual processing may rely heavily on both. In general, we also assume that language provides a powerful system for indexing simulations, and for manipulating simulations in language and thought.

1.2.4. *Caveats.* Our accounts of the “linguistic system” and the “simulation system” include simplifications that require qualification. First, we do not assume that these systems are modular, given that each is distributed throughout the brain, and that their component processes typically play roles in many cognitive activities. Second, we do not assume that each system takes the same rigid form across situations but instead is dynamical, drawing on different configurations of processes in different situations (Barsalou et al. 2007). Third, when referring to the “linguistic system,” we mean the system that processes linguistic forms, not the system that processes linguistic meaning. Clearly, meaning is a central part of language. Because we contrast linguistic forms and linguistic meaning in this article, however, we will use the “linguistic system” for the former and the “simulation system” for the latter.

1.3. *Paradigm, response coding, and predictions*

1.3.1. *Paradigm.* In two experiments, participants produced responses for similar sets of diverse cue words that referred not only to concrete concepts, but to other kinds of concepts as well (e.g. events, properties). In Experiment 1, participants generated word associations to each cue word for about 3 seconds. In Experiment 2, participants generated properties typically true of the concept associated with each cue word for 15 seconds. Including these two tasks allowed us to test a variety of predictions associated with LASS theory, as described later (1.3.4).

1.3.2. *Response coding.* Across participants, all responses to the same cue word were merged into a single master list, with minor lexical variants combined into a single response (e.g. “flower” and “flowers” formed a single response to the cue word “bee”). Two judges coded all responses in each master list, exhibiting 76% agreement in Experiment 1 and 77% agreement in Experiment 2.

A hierarchical coding scheme—presented in Appendix A—was applied sequentially to code responses. If a response was linguistically related to the cue, it was coded as a linguistically-related response. Consideration of other possible coding categories proceeded no further.³ For example, the response “hive” to the cue “bee” was coded as a linguistically-related response, because “bee hive” is a common compound phrase. Participants could have generated “hive” in response to “bee” after “bee” activated the compound linguistic form, “bee hive,” which in turn produced “hive” as a response. Possible linguistic responses included forward compound continuations (e.g. “bee” → “hive”), backward compound continuations (e.g. “bee” → “honey” from “honey bee”), synonyms (e.g., “car” → “automobile”), antonyms (e.g. “good” → “bad”), root similarity (e.g. “self” → “selfish”), and sound similarity (e.g. “bumpy” → “lumpy”).⁴ In each case, some type of linguistic relation could have related the cue word and response.

If a response did not fall into a linguistic response category, it was then evaluated for being a taxonomic response to the cue (e.g. “dog” → “animal”). If a response was taxonomically related, it was coded as a taxonomically-related response. Consideration of other possible coding categories proceeded no further.⁵ Taxonomic responses included superordinate categories (e.g. “dog” → “animal”), subordinate categories (e.g. “dog” → “terrier”), and coordinate categories (e.g. “dog” → “cat”). We realize that taxonomic relations are often included in accounts of lexical semantics, and thus could have been coded as linguistically related. Because psychologists, however, typically view taxonomic relations as conceptual, we distinguished linguistic and taxonomic responses. Later discussion and findings address this issue further.

If a response did not fall into a linguistic or taxonomic coding category, it was coded as an object or situation descriptor. Every valid response that was

3. This policy was followed because LASS theory predicts that linguistic responses should tend to be produced before non-linguistic responses (1.3.4). To test this hypothesis, any response with a linguistic relation was automatically coded as linguistic, even if it belonged to a non-linguistic coding category. See 1.3.3 for further discussion.

4. The syntax of the examples shown here is “cue” → “response”.

5. This policy was followed to test the hypothesis that taxonomic responses tend to come, at least in part, from the linguistic system and should therefore tend to be produced before object-situation responses (1.3.4). To test this hypothesis, any response with a taxonomic relation was automatically coded as taxonomic, even if it could also be coded as an object-situation descriptor. See 1.3.3 for further discussion.

not a linguistic or taxonomic response always described either a property or a situational associate of the cue concept. A property described some aspect of the cue object, whereas a situational associate was a concept that could co-occur with the cue concept in a situation. For example, “bee” produced *bee* properties (e.g. “wings”) and situational associates (e.g. “flowers”). Similarly, “golf” produced *golf* properties (e.g. “boring”) and situational associates (e.g. “sunshine”).

Finally, if a response did not fit into any of the valid response categories above, it was coded N for none. Very few responses were coded this way, with nearly all fitting into the other 10 categories.

1.3.3. *The probabilistic nature of the coding scheme.* As just described, the response “hive” to the cue “bee” could result from “bee” activating the compound linguistic form, “bee hive,” in the linguistic system, which in turn produces “hive” as a response. Importantly, however, “hive” could also result from describing a simulation of a situation containing a bee and a hive. Although this is possible, and probably occurred to some extent, we assume that this possibility is statistically less likely than “hive” originating in the linguistic system. Most importantly, we assume that linguistic responses should be *statistically more likely* to originate from linguistic processing than from simulation. As a result, linguistic responses should tend to occur early in participants’ protocols than non-linguistic responses, statistically speaking, given our assumption that the linguistic system produces responses faster than the simulation system.

The same logic applies to object-situation responses. For example, the response “flower” to the cue “bee” could result from “bee” activating a situated simulation of a bee on a flower, which in turn produces “flower” as a response. Importantly, however, “flower” could also result from an association to “bee” in the linguistic system. Most importantly, however, we assume that object-situation responses should be *statistically more likely* to result from describing simulations than to result from linguistic retrieval. Because these responses are not related linguistically to the cue, they should be less likely to originate in the linguistic system than responses that exhibit linguistic relations, statistically speaking. As a result, object-situation responses should tend to occur relatively late in participants’ protocols, given our assumption that simulations produce responses more slowly than linguistic forms.

1.3.4. *Predictions.* LASS theory predicts that the linguistic system and the simulation system should both contribute to responses in the word association task (Experiment 1) and in the property generation task (Experiment 2). In each experiment, responses produced initially should tend to originate in the linguistic system, whereas responses produced later should tend to originate in

the simulation system (Figure 1). Thus the key prediction in both experiments is that linguistically-related responses should tend to be produced earlier than object-situation responses, statistically speaking.

Second, we predicted that the linguistic system would contribute more responses for word association in Experiment 1 than for property generation in Experiment 2. In both tasks—even word association—we assume that the simulation system should produce some responses, given that this system becomes active quickly (Figure 1). Because the production period was longer for property generation than for word association, however, the simulation system should have greater opportunity to produce responses over the longer production period for property generation.

Finally, the predictions for taxonomic responses are less clear than the predictions for linguistic and object-situation responses. On the one hand, taxonomic categories are generally viewed as residing in the conceptual system (e.g. Smith 1978; Murphy 2002). On the other hand, people memorize phrases for taxonomic relations during childhood, such as “a dog is an animal.” Thus, taxonomic responses could result from retrieving linguistic forms. Furthermore, it is not clear how taxonomic categories are realized in simulations, especially superordinates. How is the superordinate *animal* evident in a situated simulation of a dog? *Animal* is not a concrete property of a dog that is simulated, nor is it a thematic associate that co-occurs with dogs in situations. These observations suggest that superordinate categories such as “animal” may originate in the linguistic system. Conversely, coordinates and subordinates may often occur as situational associates in situated simulations (e.g. a simulation of a dog as a collie, or of a dog chasing a cat). Experiment 2 offers support for these hypotheses.

2. Experiment 1

Participants produced word associates to a total of 64 cue words verbally while being recorded. To minimize carry-over effects from one trial to another, a given participant only produced word associates to 16 of the 64 concepts. By only performing 16 trials, and by keeping the cue words on these trials as different as possible, the likelihood of response carry-over was minimized.

Although participants were asked to perform word association, we predicted that their later responses on a trial would come from the simulation system. Assuming that the simulation system becomes engaged quickly after reading a cue, some responses could result from describing simulations, although these responses should generally occur later than responses from the linguistic system.

To minimize responses from the simulation system, the experimenter terminated each trial as soon as the participant paused during the response period

(typically after a few seconds), thereby preventing extended use of simulation. In contrast, participants in Experiment 2 were required to produce responses for a full 15 seconds. As a result of producing responses for less time, Experiment 1 participants should produce fewer responses from the simulation system than Experiment 2 participants. A comparison of object-simulation responses across Experiments 1 and 2 will assess this prediction later.

2.1. *Method*

2.1.1. *Design and participants.* Participants produced word associates to a total of 64 cue words, with a given participant only producing word associates to 16. Within the 64 cue words, 8 belonged to each of 8 concept types established before the experiment. Following data collection, responses to the cue words were coded by the experimenters into one of ten specific coding categories that belonged to three general coding categories (Appendix A). Thus, concept, concept-type, specific coding category, and general coding category were the four variables that structured the experiment.

Participants were 160 members of the Emory community who received \$2 for completing the experiment. Twenty participants received one of eight presentation lists (4 versions \times 2 orders), with 40 participants producing word associates for a given cue word.

2.1.2. *Materials.* The 64 cue words, shown in Appendix B, were drawn from the Nelson et al. (1999) word association norms. To make the cue words as diverse as possible, they were drawn from eight conceptual types (eight cue words per type): forward continuation, backward continuation, synonym, antonym, taxonomic category, semantic field, stereotypical object property, and brand. Each conceptual type was defined by one or more dominant associations to the cue word in the Nelson et al. norms. For example, forward continuation cues were all characterized by having a dominant word associate that was a forward continuation, namely, a word that typically follows the cue in a compound phrase (e.g. for the cue “taxi,” “cab” is a dominant response that follows it in “taxi cab”). Conversely, backward continuation cues were all characterized by having a dominant word associate that was a backward continuation, namely, a word that typically precedes the cue in a compound phrase (e.g. for the cue “muffin,” “blueberry” is a dominant response that precedes it in “blueberry muffin”). Synonym cues were all characterized by having a dominant word associate with a very similar meaning (e.g. for the cue “teacher,” the response “instructor”). Antonym cues were all characterized by having a dominant word associate with the opposite meaning (e.g. for the cue “heavy,” the response “light”). Taxonomic cues were all characterized by having a dominant word associate that was a related category (e.g. for the cue “car,” the

response “vehicle”). Semantic field cues were all characterized by having a dominant word associate that was a category from the same semantic field (e.g. for the cue “winter”, the response “summer”). Stereotypical property cues were all characterized by having a dominant word associate that was an object stereotypically containing the cue property (e.g. for the cue “wings”, the response “bird”). Brand cues were all characterized by having a dominant word associate that was a related brand (e.g. for the cue “Crest”, the response “Colgate”).

Four presentation lists were constructed that each contained 16 of the 64 cue words. In each list, two words were drawn pseudo-randomly from the eight words for each concept type, such that concept type was represented equally in every list. The 16 words in a list were adjusted slightly so that minimal semantic relatedness existed between them, thereby minimizing the chances that responses generated for one cue word would carry over to the responses for a later cue word. Within each presentation list, the words were ordered quasi-randomly, with the constraint that no two adjacent words belonged to the same conceptual type. Words were slightly reordered so as to further minimize carry-over effects. A second version of each presentation list was constructed that simply reversed the quasi-random order of the cue words in the original list, thereby creating eight total lists.

2.1.3. *Procedure.* An experimenter approached potential participants who were alone and quiet in a relatively calm campus location. The experimenter asked whether the person was interested in performing a 10 minutes psychology experiment for \$2. If the participant was clearly willing, he or she was asked for informed consent, otherwise the experimenter departed.

The experimenter then read instructions to the participant, stating first that the experiment used a simple word association task to study people’s knowledge about words. Participants were further told that the experimenter would say a word, that the subject should note the first words that came to mind, and that the participant should then say these words out loud. The instructions stressed that there were no correct answers, that the experiment would be most successful if participants responded naturally and spontaneously with whatever words came to mind initially, and that the words produced did not need to describe a coherent definition or situation. The instructions also stressed that there was no need to say any further words that might come to mind after those that came to mind immediately.

The first two trials prior to the 16 critical trials were practice, although the participant was not aware that practice was being performed. The two cue words for the practice trials were drawn from two different conceptual types in a different list. The practice cues differed for each of the four list versions, selected to minimize the possibility of carry-over effects to later trials. The experimenter recorded the 18 trials on a hand-held digital audio recorder.

On each trial, the experimenter first stated, “For the following word, what other words come to mind immediately” and then stated the cue word. The participant responded verbally with word associates. As soon as the participant paused, the experimenter ended the trial with “That’s fine,” and proceeded to the next trial.

2.2. Results

Participants produced an average of 1.74 responses for each word cue over an average production period of 2.98 seconds (the end of the production period was defined as the onset of silence after the final verbal response). As described earlier, these responses occurred before the participants paused, thereby limiting use of the simulation system in producing responses.

Table 1 presents the proportions of responses that fell into each of the specific and general response categories. The results for the specific response categories validate the assignment of concepts to concept types. The largest proportion of forward continuation responses (CF) occurred for concepts in the forward compound continuation condition, whereas the largest proportion of backward continuation responses (CB) occurred for concepts in the backward compound continuation condition. Similarly, the largest proportion of synonym responses (SN) occurred for synonym concepts, whereas the largest proportion of antonym responses (AN) occurred for antonym concepts. Finally, the largest proportions of taxonomic responses (DH, DL, DS) occurred for semantic field and taxonomic concepts. Brand concepts also produced taxonomic responses frequently, given that brands often activated the superordinate categories (DH) to which they belong.

This pattern of results indicates that Experiment 1 tapped into word association patterns reported in previous research on word association (Nelson et al. 1999). Word association appears to be a stable process that can be replicated in different samples over the course of a decade. Our procedure tapped into this process, given that it produced classic patterns of word association.

The results for the general response categories on the right side of Table 1 indicate that responses were roughly balanced across the three categories, with 34% linguistic responses, 27% taxonomic responses, and 38% object-situation responses. The presence of 38% object-situation responses suggests that responses from the simulation system contributed to production, even though participants were asked to produce word associates in a short time period. Further interpretation of this finding will be presented later.

Table 2 presents the central results for average output position of responses. This analysis was performed on the 1,227 unique responses generated to the 64 word cues, where a unique response was often generated by multiple participants, and sometimes took slightly different forms (e.g. “flower” and “flowers”

were combined into a single response). The dependent measure was the median output position of the response across the participants who produced it, where 1 indicated the first response to a cue word, 2 the second response, 3 the third response, and so forth.⁶ The median output positions of the unique responses were entered into a concept type \times general response code ANOVA, using concepts as the random factor.⁷

As can be seen on the right side of Table 2, the results for the general response categories support LASS theory. Linguistic responses were produced earliest (1.61), followed by taxonomic responses (2.03), and then object-situation responses (2.47). The omnibus effect of general response category on output position was significant, $F(2, 1104) = 27.70$, $MSE = 1.92$, $p < .001$. Planned comparisons further found that linguistic responses were faster than taxonomic responses, $F(2, 1104) = 16.95$, $MSE = 1.92$, $p < .001$, and that taxonomic responses were faster than object-situation responses, $F(2, 1104) = 19.51$, $MSE = 1.92$, $p < .001$. As the right side of Table 2 further illustrates, linguistic responses had faster averages than taxonomic responses for seven of the eight concept types. In turn, taxonomic responses had faster averages than object-situation responses for seven of the eight concept types. The consistency of this pattern was reflected in the absence of an omnibus interaction between concept type and general response category, $F(14, 1104) = 1.07$, $MSE = 1.92$, *ns*.

Finally, the results for specific response categories in Table 2 confirm the pattern of results for the general response categories. As can be seen, the specific linguistic categories tended to be faster than the specific taxonomic categories, which tend to be faster than the object-situation category. The one exception was root similarity responses, which tended to have late output positions. As Table 1 illustrates, root similarity responses only occurred 1% of the time, such that the value for mean output position is likely to be unreliable.

2.3. Discussion

Consistent with LASS theory, linguistic responses to word association cues tended to occur before object-situation responses. These results are consistent with the proposal that responses originate in different systems for linguistic

6. Analyses performed on the mean time to produce a response showed virtually the same results.

7. In the inferential tests to follow, concepts constituted the random factor for several reasons. First, each participant only produced responses for 16 of the 64 concepts. Second, Raaijmakers (2003) shows that when participants receive different lists of items from the same population, taking both participants and items (i.e. concepts) into account is not necessary, given that both sources of variability enter into a single F test.

Table 1. Proportion of responses by specific response category, general response category, and concept type from Experiment 1

Concept Type	Specific Response Category										General Response Category			
	CF	CB	SS	RS	SN	AN	DH	DL	DS	OS	Ling	Tax	OS	
CC Forward	.44	.03	.01	.03	.12	.00	.04	.01	.03	.31	.62	0.08	0.31	
CC Backward	.08	.48	.00	.00	.08	.00	.03	.01	.04	.32	0.60	0.09	0.32	
Synonym	.03	.01	.00	.04	.23	.03	.09	.02	.09	.45	0.34	0.19	0.45	
Antonym	.10	.02	.01	.01	.08	.19	.00	.01	.17	.41	0.43	0.18	0.41	
Taxonomic	.08	.07	.00	.00	.00	.00	.23	.01	.26	.36	0.15	0.26	0.36	
Semantic Field	.06	.01	.00	.01	.00	.00	.10	.05	.45	.33	0.08	0.59	0.33	
Property	.08	.13	.00	.01	.08	.08	.03	.01	.11	.47	0.39	0.15	0.47	
Brand	.14	.01	.00	.01	.01	.00	.32	.00	.09	.42	0.17	0.41	0.42	
Overall Mean	.13	.09	.00	.01	.08	.04	.10	.01	.15	.38	0.34	0.27	0.38	

Notes: CC = compound continuation, CF = compound continuation forward, CB = compound continuation backward, SS = sound similarity, RS = root similarity, SN = synonym, AN = antonym, DH = domain higher-level category, DL = domain lower-level category, DS = domain same-level category, OS = object or situation descriptor, Ling = linguistic, Tax = Taxonomic.

Table 2. Average output position during word association by specific response category, general response category, and concept type from Experiment 1.

Concept Type	Specific Response Category										General Response Category		
	CF	CB	SS	RS	SN	AN	DH	DL	DS	OS	Ling	Tax	OS
CC Forward	1.59	1.80	1.00	2.00	1.59		3.08	1.50	1.60	2.08	1.61	2.27	2.08
CC Backward	1.75	1.67			2.50		2.29	2.50	1.57	2.55	1.75	2.00	2.55
Synonym	1.67	1.00		3.00	1.48	1.00	2.11	1.67	1.91	2.33	1.54	1.96	2.33
Antonym	1.85	1.00	1.00	1.00	1.78	1.25		3.67	1.78	2.03	1.56	1.98	2.03
Taxonomic	1.44	1.00		3.00	1.00		1.77	1.00	1.78	2.42	1.38	1.74	2.42
Semantic Field	2.75	2.00		3.00			2.32	1.72	2.69	2.96	2.64	2.45	2.96
Property	1.29	1.30	1.00	1.50	1.23	1.25	1.67	1.00	1.60	2.23	1.27	1.58	2.23
Brand	1.90	1.50		1.00	2.00	1.00	2.06		2.27	3.21	1.76	2.16	3.21
Overall Mean	1.71	1.53	1.53	2.00	1.58	1.21	2.10	1.89	2.02	2.47	1.61	2.03	2.47

Notes: Empty cells contained no responses. CC = compound continuation, CF = compound continuation forward, CB = compound continuation backward, SS = sound similarity, RS = root similarity, SN = synonym, AN = antonym, DH = domain higher-level category, DL = domain lower-level category, DS = domain same-level category, OS = object or situation descriptor, Ling = linguistic, Tax = Taxonomic.

forms and simulations, with responses from the linguistic form system tending to originate first.

Interestingly, taxonomic responses tended to be intermediate in output position. One possibility is that they originate equally often in the language and simulation systems, in contrast to linguistic responses that originate most often during linguistic processing and to object-situation responses that originate most often during simulation. On some trials, superordinate, subordinate, and coordinate concepts are produced as word associations; on others, they are produced while construing and describing a situated simulation (see Barsalou 2003a, 2005a, 2008b for an account of how such construal might occur). Experiment 2 offers further insight into the production of taxonomic associates.

Although participants were instructed to produce word associates, they also appeared to draw on the simulation system to produce some responses. Specifically, 38% of their responses did not reflect linguistic relations or taxonomic relations to the word cues. Of course, it is possible that these responses nevertheless originated in the word association system (e.g. the cue words activated these words without simulations becoming active; see 1.3.3). The fact that object-situation responses were significantly slower than linguistic responses, however, suggests that many object-situations responses originated in a different system. If object-situation responses had exhibited the same average output position as linguistic responses, this would have suggested that both types of responses originated in a common system. The fact that linguistic responses were substantially faster than object-situation responses suggests otherwise. Findings from the next experiment further support this conclusion, as does the Simmons et al. (2008) finding that later responses to word cues activate simulation areas in the brain (also see Kan et al. 2003).

3. Experiment 2

Experiment 2 assessed predictions of LASS theory in the property generation task. On each trial, participants received a cue word and produced typical properties of the underlying concept verbally for 15 seconds while being recorded. Of primary interest was whether linguistically-related responses would again tend to precede object-situation responses. Of secondary interest was whether a higher proportion of responses from the simulation system would be observed in Experiment 2 than in Experiment 1. Because participants here had to produce properties for a longer duration than did participants in Experiment 1 (15 seconds vs. approximately 3 seconds), the simulation system should have more opportunity to produce simulation responses (Figure 1).

Finally, data from Experiment 1 allowed us to look at the effect of an additional variable in Experiment 2. Each concept in Experiment 1 was normed for the percentage of linguistic responses that it produced, with this percentage

ranging from 0% to 96%. Whereas some cue words never produced a linguistic response in Experiment 1, some produced mostly linguistic responses. The 60 cue words in Experiment 2 were divided into three groups of 20 that produced high (59%), medium (33%), and low (6%) levels of linguistic responses on the average in Experiment 1. Assessing this variable allowed us to further assess the role of word association during property generation, and to further assess the relative contributions of the linguistic and simulation systems.

3.1. *Method*

3.1.1. *Design and participants.* Participants produced properties of the concepts named by the 60 cue words, with a given participant only producing word associates to 30 of the cues. Within the 60 cue words, subsets belonged to different concept types established before the experiment. Following data collection, responses to the cue words were coded by the experimenters into one of ten specific coding categories that belonged to three general coding categories. Thus, concept, concept-type, specific coding category, and general coding category were the four variables that structured the experiment.

Participants were 12 members of the Emory community who received \$10 for completing the experiment (none participated in Experiment 1). Two participants received one of six presentation lists (2 versions \times 3 orders), with six participants producing word associates for a given cue word. Only twelve participants were used so that the number would be comparable to the companion neuroimaging experiment, thereby facilitating comparison. As will be seen, robust effects occurred in this sample.

3.1.2. *Materials.* The cue words from Experiment 1 were also used here. Because of design constraints in the companion fMRI experiment (Simmons et al. 2008), only 60 of the original 64 cue words were used. The four words dropped (indicated in Appendix B) were selected so as to optimize the manipulation of linguistic associatedness in this experiment (i.e. dropping these four words minimized differences between levels of linguistic associativeness across different list versions).

The companion neuroimaging experiment required that the 60 cue words be divided into two lists of 30 words each. A given participant in the neuroimaging experiment received one list in a first neuroimaging session that assessed property generation, and then received the other list in a second neuroimaging session that assessed word association and situation generation, with the assignment of lists to sessions counter-balanced across participants. In the experiment here, each participant received only one of the 30-word lists from the neuroimaging experiment in a single session. By only presenting par-

ticipants here with one of these lists, we replicated the behavioral task that neuroimaging participants performed in their initial critical scanning session.⁸

Assignments of the 30 words in each list reflected the following three constraints. First, the 30 cue words had to be balanced in terms of the number drawn from each conceptual type (three or four per type). Second, the 30 cue words had to be balanced in terms of their linguistic associativeness. Based on the results of Experiment 1, the 30 cues in each list were grouped into three sets of ten word cues that differed in terms of whether they had produced a high, medium, or low percentage of linguistic responses. Thus, the two lists were equated in terms of both conceptual types and the three levels of linguistic associativeness.

Each 30 word list was then structured exactly as in the companion neuroimaging experiment. Because the neuroimaging experiment had five critical runs that contained six cue words each, the lists here also contained five critical sub-lists each containing six cue words. Because each run in the neuroimaging experiment contained three blocks of two property generation trials, the lists here mirrored this structure. Because the six property generation trials in each run used two word cues that were high in linguistic associativeness, two that were medium in linguistic associativeness, and two that were low in linguistic associativeness, the lists here contained the same distribution of cues. Because the neuroimaging experiment had blocks of six lexical decision trials (three words and three non-words) that followed each block of two property generation trials, the lists here also contained lexical decision blocks. Because the lexical decision trials were simply used as fillers in the neuroimaging experiment, they were also treated as fillers here. Words in the lexical decision task were unrelated to the cue words.

In summary, the materials contained five critical sub-lists. Each sub-list contained, in sequence, two property generation trials, six lexical decision trials, two property generation trials, six lexical decision trials, two property generation trials, and six lexical decision trials, where each block of two property generation trials used two word cues that were high, medium, or low in linguistic associativeness. Three different versions of each 30-word list were constructed that counter-balanced the order of high, medium, and low

8. We document relations between the current experiment and the neuroimaging experiment in detail for several reasons. First, some aspects of the current experiment that may seem somewhat arbitrary make more sense when the constraints that existed on designing the companion neuroimaging experiment are noted. Second, by noting the close correspondence between the behavioral and neuroimaging experiments, we are justified later in claiming that the neuroimaging results inform the interpretation of the behavioral results, and vice versa. Third, providing detail about the correspondence between the two experiments will assist those readers who want to examine relations between the two experiments closely.

linguistic associativeness across blocks and runs. A practice list having the same structure as the five critical lists was also constructed that had comparable property generation and lexical decision trials, using materials unrelated to the critical list.

3.1.3. *Procedure.* The experimenter first described the property generation and lexical decision tasks to the participant. For the property generation task, the experimenter asked participants to think of the characteristics typically true of the concept named by a cue word (e.g. what characteristics are typically true of *taxi*) and to produce them verbally. Participants were instructed that there were no correct answers, and that they should produce whatever characteristics came to mind, as they came to mind spontaneously.

Based on a variety of other experiments, the duration of the property generation period was set at 15 seconds, which was the same duration used in the companion neuroimaging experiment. Unlike word association, which typically ends in a few seconds, property generation typically extends to about 15 seconds before participants begin pausing extensively and appear to shift strategy.

The behavioral sequence of events in the experiment here was identical to the behavioral sequence in the companion neuroimaging experiment. Six runs were performed, with the first being practice, and the next five being critical. Each run contained three alternating cycles of a property generation block for two cue words followed by a lexical decision block for six strings.

Prior to each property generation block, “Properties?” appeared on the computer screen, instructing participants that they were to generate properties for the next two cue words. Each cue word then appeared and the participant began producing properties verbally for 15 seconds.⁹ After the second 15 seconds period ended, “Lexical?” appeared on the screen, instructing participants that they were to judge whether next six strings were words or non-words, with each trial lasting 5 seconds. Following the three alternating cycles of property generation and lexical decision, a short break occurred, comparable to the break between runs in the companion neuroimaging experiment. The practice

9. Within the 15 seconds generation period, the cue word was presented three times in three 5 seconds cycles. In each cycle, the cue was presented for 4850 ms, followed by a blank screen for 150 ms. This presentation method was used in the companion neuroimaging experiment to equate visual stimulation between the property generation condition and the word association condition. In the word association condition, the generation period was only 5 seconds, with the cue word being on the screen for the first 4850 ms, followed by a 150 ms blank screen. Presenting the same cue word three times in a 15 seconds property generation condition was comparable—in terms of visual stimulation—to presenting three different cue words on three consecutive 5 seconds word association trials. To replicate the behavioral procedure of the companion neuroimaging experiment as closely as possible, the same procedure was used here.

run segued into the five critical runs, with participants not aware of the distinction. A digital audio recorder captured participants' verbal responses for both the property generation and lexical decision tasks.

3.2. Results

Participants produced an average of 5.87 responses per cued word over the 15 seconds property generation period. As expected, participants here produced many more responses than did participants in Experiment 1, who produced an average of 1.74 responses in 2.98 seconds.

All statistical tests in Experiment 2 were performed on both participants and items (concepts) as random factors.¹⁰ When the random factor was participants, the median value across all relevant concepts (cue words) was entered into each cell of the analysis design; analogously, when the random factor was concepts, the median value across all relevant subjects was entered into each cell.¹¹ F statistics that resulted from participant analyses are labeled F_P , whereas F statistics that resulted from concept analyses are labeled F_C . Analyses on proportions were performed on arcsine transformations of those data (Winer 1971).

Table 3 presents the proportions of responses that fell into each of the specific and general response categories. The results for the specific response categories again validate the assignment of concepts to concept types. Interestingly, however, these patterns are not as clear-cut as in Experiment 1, probably because the longer production interval allowed more responses from the simulation process to mask characteristic word association patterns from the linguistic system. As in Experiment 1, the largest proportion of forward continuation responses (CF) occurred for concepts in the forward compound continuation condition, whereas the largest proportion of backward continuation responses (CB) occurred for concepts in the backward compound continuation condition. Notably, however, the proportions here were much smaller than in Experiment 1. The largest proportion of synonym responses (SN) occurred for antonym concepts, although synonym concepts produced the next highest proportion. Again, antonyms produced the largest proportion of antonym responses (AN), but to a much smaller degree than in Experiment 1. Finally, the largest proportions of taxonomic responses (DH, DL, DS) again

10. Similar to Experiment 1, it was technically not necessary to include both participants and concepts as random factors, given that participants received different lists of items from the same population, such that both sources of variability entered into a single F test (Raaijmakers 2003). Nevertheless, both random factors were included so that the variability of each could be assessed.

11. Analyses performed on means across responses in the same cell of the design, instead of medians, showed virtually the same results.

Table 3. Proportion of responses by specific response category, general response category, and concept type from Experiment 2.

Concept Type	Specific Response Category										General Response Category			
	CF	CB	SS	RS	SN	AN	DH	DL	DS	OS	Ling	Tax	OS	
CC Forward	.27	.01	.00	.01	.07	.00	.08	.02	.02	.52	.36	.12	.52	
CC Backward	.06	.19	.00	.02	.05	.00	.05	.06	.03	.53	.33	.14	.53	
Synonym	.04	.00	.00	.03	.19	.00	.09	.05	.05	.55	.26	.19	.55	
Antonym	.06	.02	.01	.00	.27	.04	.07	.01	.05	.47	.40	.13	.47	
Taxonomic	.06	.05	.00	.00	.01	.00	.17	.05	.05	.61	.12	.27	.61	
Semantic Field	.10	.01	.01	.02	.00	.00	.15	.06	.12	.52	.14	.34	.52	
Property	.06	.07	.00	.00	.11	.02	.14	.00	.11	.49	.26	.25	.49	
Brand	.13	.02	.00	.01	.01	.00	.24	.03	.00	.57	.16	.27	.57	
Overall Mean	.10	.04	.00	.01	.09	.01	.12	.04	.05	.54	.24	.22	.54	

Notes: CC = compound continuation, CF = compound continuation forward, CB = compound continuation backward, SS = sound similarity, RS = root similarity, SN = synonym, AN = antonym, DH = domain higher-level category, DL = domain lower-level category, DS = domain same-level category, OS = object or situation descriptor, Ling = linguistic, Tax = Taxonomic.

Table 4. Proportion of responses by linguistic association level and general response category from Experiment 2.

Linguistic Associativeness	General Response Category		
	Linguistic	Taxonomic	Object & Situation
High	.33	.14	.53
Medium	.27	.20	.53
Low	.12	.33	.55

occurred for semantic field, taxonomic, and brand concepts. Again, however, these advantages were much smaller than in Experiment 1.

This pattern of results indicates that Experiment 2 tapped into characteristic word association patterns reported in previous research on word association (Nelson et al. 1999). More significantly, this pattern indicates that word association is part of the property generation process. Contrary to standard accounts of property generation, properties are not produced solely by accessing conceptual representations. Instead, a significant proportion of responses on this task appears to originate in a word association process.

The results for the general response categories on the right side of Table 3 indicate that the global pattern of responding differed considerably between Experiments 1 and 2. Most importantly, the proportion of object-situation responses increased from 38% in Experiment 1 to 54% in Experiment 2 (a 42% increase). Conversely, the proportions of taxonomic and linguistic responses decreased across experiments (34% and 27% in Experiment 1 versus 24% and 22% in Experiment 2). As predicted, the longer production period in Experiment 2 appeared to increase responses from the simulation system. Object-situation responses here were much more frequent than linguistic responses ($F_p(1,22) = 30.36$, $MSE = .461$ arcsine units, $p < .001$, $F_C(1,104) = 63.34$, $MSE = .194$ arcsine units, $p < .001$) and were also much more frequent than taxonomic concepts ($F_p(1,22) = 48.15$, $MSE = .461$ arcsine units, $p < .001$, $F_C(1,104) = 73.62$, $MSE = .194$ arcsine units, $p < .001$).

Table 4 presents results from the manipulation of linguistic associativeness. As described earlier, the 60 cue words were normed for the proportion of linguistic responses that they produced in Experiment 1. Table 4 illustrates that the proportion of linguistic responses in Experiment 2 decreased significantly as linguistic association decreased in Experiment 1, $F_p(1,44) = 37.91$, $MSE = .038$ arcsine units, $p < .001$, $F_C(1,114) = 21.79$, $MSE = .188$ arcsine units, $p < .001$. This effect further implicates word association as a central process in property generation. Because the proportion of language-based responses produced on a word association task strongly predicts the proportion of linguistic

responses here, word association appears to be a central component of property generation.

Interestingly, linguistic and taxonomic responses were inversely related in Experiment 2 across the three levels of linguistic associativeness. As we just saw, the proportion of linguistic responses decreased as linguistic associativeness decreased. In contrast, the proportion of taxonomic responses increased as linguistic associativeness decreased, $F_p(1,44) = 34.88$, $MSE = .038$ arcsine units, $p < .001$, $F_C(1,114) = 13.30$, $MSE = .188$ arcsine units, $p < .001$. Perhaps most interestingly, the proportion of object-situation responses remained almost exactly constant across the three levels of linguistic associativeness, $F_p(1,44) = .06$, $MSE = .038$ arcsine units, ns , $F_C(1,114) = .09$, $MSE = .188$ arcsine units, ns . The overall pattern in Table 4 suggests that linguistic and taxonomic responses tap a single process responsible for producing linguistic and taxonomic responses, whose output remains relatively constant across large changes in the linguistic associativeness of concepts. Conversely, this pattern further suggests that object-simulation responses tap a second process, whose output also remains relatively constant across large changes in linguistic associativeness. More generally, this pattern supports LASS theory's central assumption that two processes underlie conceptual tasks such as property generation, with the linguistic process being largely responsible for linguistic and taxonomic responses.

Table 5 presents the central results for the output position of responses, where 1 indicates the first response to a cue word, 2 the second response, 3 the third response, and so forth. As can be seen on the right side of Table 5, the results for the general response categories support LASS theory. Linguistic responses were produced fastest (3.30), taxonomic responses were produced just as fast (3.33) and object-situation responses were produced more slowly (4.25). The omnibus effect of general response category on output position was significant, $F_p(2,22) = 11.57$, $MSE = 2.35$, $p < .001$, $F_C(2,104) = 15.14$, $MSE = 1.16$, $p < .001$. Planned comparisons found further that linguistic and taxonomic responses did not differ, $F_p(1,22) = .02$, $MSE = 2.35$, ns , $F_C(1,104) = .02$, $MSE = 1.16$, ns . Linguistic responses, however, were faster than object-situation responses, $F_p(1,22) = 18.43$, $MSE = 2.35$, $p < .001$, $F_C(1,104) = 23.34$, $MSE = 1.16$, $p < .001$; taxonomic responses were also faster than object-situation responses, $F_p(1,22) = 17.29$, $MSE = 2.35$, $p < .001$, $F_C(1,104) = 21.89$, $MSE = 1.16$, $p < .001$. As the right side of Table 5 illustrates further, linguistic responses had faster averages than object-situation responses for all eight concept types, and taxonomic responses had faster averages than object-situation responses for six of the eight concept types. Unlike Experiment 1, an interaction between general response category and concept type indicated that concept type modulated the average output positions of the general response categories, $F_p(14,154) = 4.21$, $MSE = 1.13$, $p < .001$,

Table 5. Average output position during word association by specific response category, general response category, and concept type from Experiment 2

Concept Type	Specific Response Category										General Response Category		
	CF	CB	SS	RS	SN	AN	DH	DL	DS	OS	Ling	Tax	OS
CC Forward	3.58	2.00		3.67	3.00		2.76	4.60	5.75	4.35	3.45	3.58	4.35
CC Backward	4.00	3.31		4.00	2.69		3.43	5.24	6.30	4.23	3.38	4.88	4.23
Synonym	4.78	5.00	1.50		2.67	3.00	4.05	2.92	3.36	3.76	2.92	3.55	3.76
Antonym	3.23	2.25	1.00		2.85	2.22	2.25	3.00	3.10	3.91	2.79	2.61	3.91
Taxonomic	4.17	4.18		1.00	1.50		2.50	5.12	3.00	4.54	3.95	3.09	4.54
Semantic Field	4.11	2.00	5.00	1.40	2.00		2.22	4.47	3.67	4.05	3.63	3.16	4.05
Property	2.38	3.06			3.57	2.80	3.17	2.00	4.54	3.72	3.09	3.74	3.72
Brand	3.53	4.80		6.50	4.50		2.65	3.22	5.00	4.93	3.80	2.74	4.93
Overall Mean	3.67	3.42	3.67	3.83	2.92	2.47	2.74	4.34	4.04	4.25	3.30	3.33	4.25

Notes: The value in each non-empty cell of this table is the mean of all responses for that cell in the design out of the 2,114 total responses generated across participants and concepts (i.e., these values are not medians across subjects or concepts as in the ANOVAs). Empty cells contained no responses. CC = compound continuation, CF = compound continuation forward, CB = compound continuation backward, SS = sound similarity, RS = root similarity, SN = synonym, AN = antonym, DH = domain higher-level category, DL = domain lower-level category, DS = domain same-level category, OS = object or situation descriptor, Ling = linguistic, Tax = Taxonomic.

$F_C(14,104) = 3.83$, $MSE = 1.16$, $p < .001$. This difference between experiments suggests that greater use of simulation caused concepts to diverge in their response profiles.

As Table 5 illustrates, the three taxonomic response categories differed considerably in their response profiles. Superordinate (DH) responses (2.74) were much faster than subordinate (DL) responses (4.34) and coordinate (DS) responses (4.04). Notably, superordinate responses (2.74) were the second fastest specific category in the experiment, with only antonym (AN) responses being slightly faster (2.47). Otherwise the superordinate responses were faster than all other linguistic categories. Conversely, the subordinate (4.34) and coordinate responses (4.04) were the first and third slowest categories in the experiment, comparable to object-situation (OS) responses (4.25), which were second slowest. As suggested shortly, superordinate responses may have tended to originate in the linguistic system, whereas subordinate and coordinate responses may have tended to originate in the simulation system.

3.3. Discussion

Experiment 2, like Experiment 1, supports predictions of LASS theory. Linguistic responses tended to occur early during property generation, whereas object-situation responses tended to occur late, suggesting that linguistic and object-situations responses originated in two systems.

Taxonomic responses behaved somewhat differently in Experiment 2 than in Experiment 1. In Experiment 1, the output position of taxonomic responses was intermediate between the output positions of linguistic and object-situation responses. In Experiment 2, taxonomic responses were as fast as linguistic responses, suggesting that both tended to originate in the same system. The finding that linguistic and taxonomic responses summed to a constant proportion (Table 4) further supports this conclusion.

The results for different types of taxonomic responses qualify this conclusion somewhat. In Experiment 2, superordinate concepts (DH) were some of the fastest responses produced, suggesting that they typically originated in the linguistic system. Participants appeared to generate superordinates primarily through word association. Conversely, subordinate and coordinate concepts (DL and DS) were as slow as object-situation responses, suggesting that they originated in simulation. Subordinate responses may have resulted from simulating a specific form of the cue concept and then categorizing it (e.g. for the cue “car,” simulating a sedan, categorizing it as *sedan*, and producing “sedan” verbally). Coordinate concepts may have resulted from simulating a coordinate concept that typically co-occurs with the cue concept in the same situation (e.g. for the cue “car,” simulating a situation that includes another vehicle

likely to occur, such as *truck*, and then categorizing and naming the truck to produce a response).

Nothing in the instructions for the property generation task mentioned or implicated word association. Nevertheless, the results of Experiment 2 indicated that word association spontaneously entered into property generation. First, property generation continued to exhibit characteristic patterns of word association for different conceptual types (Table 3). Second, the linguistic associativeness of word cues obtained from Experiment 1 predicted the proportion of linguistic responses generated to cues in Experiment 2. Thus property generation is not a pure conceptual task that only utilizes conceptual representations. Instead, it draws heavily on word associations from the linguistic system.

Although word association was important for property generation, simulation was even more important, with object-situation responses being more frequent than linguistic and taxonomic responses. Furthermore, the proportion of object-situation responses was much higher in Experiment 2 (54%) than in Experiment 1 (38%). Because Experiment 2 had a longer generation period, the simulation system had more time to produce object-situation responses.

The higher percentage of simulation responses in Experiment 2 appeared to produce two further differences between Experiments 1 and 2. First, the characteristic patterns of word association for different conceptual types were weaker in Experiment 2 than in Experiment 1 (Table 3 vs. Table 1). Higher proportions of simulation responses appeared to mask characteristic word association patterns. Second, different conceptual types exhibited more variability for the average output positions of general response categories in Experiment 2 than in Experiment 1 (i.e. concept type and general response category interacted in Experiment 2 but not in Experiment 1; Table 5 vs. Table 2). Higher proportions of simulation responses appeared to create greater diversity in the types of responses produced for different conceptual types, suggesting that the greatest differences between conceptual types may reside in simulated content as opposed to word associations.

4. General discussion

LASS theory proposes that verbal responses to word cues do not result from a unitary conceptual system composed of amodal representations. Instead, verbal responses to word cues originate in two systems—the linguistic form system and the simulation system—with the linguistic form system tending to produce responses earlier than the simulation system (Figure 1). The results of Experiments 1 and 2 support this account. When participants generated word associations to word cues in Experiment 1, responses linguistically related to the cues were produced earlier on average than responses not linguistically

related. Similarly, when participants generated properties for concepts in Experiment 2, responses related linguistically to the cues were produced earlier.

Several additional findings corroborate this primary result. First, the overall proportion of linguistically related responses was lower for property generation than for word association, consistent with the prediction that word associations should be less frequent when more time is available for the slower simulation process. Second, the specific types of word associations produced to word cues in both experiments corresponded to associations reported in the Nelson et al. (1999) norms, implicating word association. Furthermore, this pattern was weaker in Experiment 2 than in Experiment 1, given that greater opportunity for simulation during property generation diluted and distorted contributions from word association. Third, the proportion of linguistic responses to the cue words in Experiment 1 during word association strongly predicted the proportion of linguistic responses to the same words in Experiment 2 during property generation. Even when participants were instructed to perform conceptual property generation, word associations entered into the responses they produced.

4.1. *Corroborating results from the companion neuroimaging experiment*

In a companion study to Experiment 2, Simmons et al. (2008) scanned participants on two occasions a week apart. During the first scanning session, participants received visual cue words for 30 concepts from Experiment 2 (e.g. “car,” “heavy,” “compute,” “guilty”). As participants read a word, they generated properties of the associated concept to themselves for 15 seconds (i.e. the property generation task). In the second scanning session a week later, the same participants performed two localizer tasks that made it possible to test LASS theory. Participants received additional concepts not seen in the first session. For some concepts, participants generated word associates for 5 seconds (e.g. for *car*, “automobile,” “vehicle”). For other concepts, participants imagined a situation that contained the concept for 15 seconds each (e.g. for *car*, a participant might imagine a car driving along a neighborhood street). Concepts were counterbalanced so that each concept occurred in the property generation, word association, situation simulation conditions, with each participant receiving a concept once in a blocked design.

Predictions for the two localizer tasks were as follows. First, the word association task should activate left hemisphere language areas, especially Broca’s area. Second, the situation simulation task should activate bilateral posterior areas typically involved in mental imagery, such as the precuneus. The localizer results supported these predictions. Areas more active for word association than for situated simulation included a large activation in left inferior frontal gyrus (Broca’s area), along with large activations in left inferior

temporal gyrus and right cerebellum. As Simmons et al. review, these areas have often been reported in research on word processing, especially word production. Conversely, areas more active for situated simulation included a large activation in the precuneus, along with a large activation in right middle temporal gyrus. An area in the right middle frontal gyrus was also active, but at a lower significance level. As Simmons et al. review, these areas are often associated with various forms of simulation that underlie mental imagery, episodic memory, and situational processing.

Activations during the property generation task the week before were of primary interest. If LASS theory is correct, then two findings should occur. First, if property generation results from word association and situated simulation, then activations during property generation should overlap with activations during the two localizer tasks. Second, activations for property generation in word association areas should tend to occur earlier than activations for property generation in situated simulation areas.

To test these hypotheses, each 15 seconds property generation block for a single word cue was divided into two smaller 7.5 seconds blocks for the early vs. late phases of property generation. Across word cues, brain areas active early during property generation were identified, as were brain areas active later. Consistent with LASS theory, property generation areas active early overlapped with areas observed for the word association localizer (left inferior frontal gyrus and right cerebellum). Conversely, property generation areas active late overlapped with areas observed for the situation localizer (precuneus and right middle temporal gyrus). As in Experiment 2 here, two systems appeared responsible for producing responses during property generation: The linguistic system and the situated simulation system, with the linguistic system producing responses earlier.

4.2. *Integrating the behavioral and neuroimaging results*

The behavioral results here and the neuroimaging results from Simmons et al. (2008) provide stronger support for LASS theory than either set of results alone. On the one hand, the behavioral results indicate that property generation responses associated with linguistic relations tend to be produced earlier than responses not associated with linguistic relations. On the other hand, the neuroimaging results indicate that property generation activations in word association areas tend to occur earlier than activations in situated simulation areas. The behavioral and neuroimaging results corroborate each other by linking complementary markers of linguistic processing with early responses (linguistic relations, word association areas), and by linking complementary markers of situated simulation with later responses (object-situation properties, situated simulation areas). Together, these findings suggest that two systems underlie

verbal responses during property generation, with the linguistic system producing responses earlier than the simulation system.

As described earlier, the results of the two experiments here also conform to previous theories and results. Theoretically, these results conform to the general outline of Paivio's (1971, 1986) Dual Code Theory and to Glaser's (1992) extension of it. Empirically, these results conform to the results reported in Chaffin (1997), Solomon and Barsalou (2004), Kan et al. (2003), Louwse (2008), and Andrews et al. (2009). Together these theories and results suggest that the linguistic system and the simulation system must both be included in accounts of conceptual processing.

4.3. *Implications for word association and property generation*

If the account as proposed by LASS theory is implicated in conceptual processing, then implications follow for collecting concept norms based on verbal responses to word cues.¹² First, different properties in the norm collected for a concept may originate in different systems. Some properties may originate consistently in the linguistic system, some may originate consistently in the simulation system, and some may originate in both. Second, properties originating in different systems may have different predictive capabilities. If a concept norm is used to predict performance on other conceptual tasks, such as categorization, inductive inference, or conceptual combination, then properties originating in different systems may have differential predictive success for these other tasks. Similarly, if a concept norm is used to predict linguistic behavior, different sources of information in the concept may have different predictive ability for different aspects of performance. Finally, when a concept norm for a social group is used to predict social perception or behavior, different sources of information may predict social interaction differently.

Furthermore, the results emerging from this literature suggest that one must be careful in drawing conclusions about the mechanisms underlying a particular linguistic task. Consider word association. Although researchers typically assume that word production mechanisms typically underlie word association, current evidence suggests that simulation contributes as well. In Wu and Barsalou (2009), the perceptual variable of occlusion consistently affected verbal responses during a word association task, indicating the presence of simulation. Similarly, Solomon and Barsalou (2004) found that simulation mechanisms affected property verification even under task conditions strongly favoring word association.

12. For implications that bear on other tasks and non-linguistic cues (e.g. pictures), see Glaser (1992), Barsalou et al. (2008) and Simmons et al. (2008).

Similarly, researchers must be careful to avoid concluding mistakenly that a conceptual task—such as property generation or property verification—reflects pure “conceptual processing.” As we have seen here, property generation includes responses from word association. Solomon and Barsalou (2004) also found that word association affected property verification under conditions favoring simulation. In both property generation and property verification, mixtures of word association and simulation contribute to performance.

Together, these findings suggest that a given task typically contains a varying mixture of underlying processes. In property generation, mixtures of linguistic form processing and situated simulation appear present to varying extents. As Experiments 1 and 2 here suggest, the mixture of these two processes varies between word association and property generation, with situated simulation being more important for property generation than for word association. Similarly, in the Solomon and Barsalou (2004) and Kan et al. (2003) experiments on property verification, different types of false trials modulated the presence of simulation during property verification. As these examples illustrate, mixtures of processes may typically be present on a task, rather than a single unitary process, and be modulated by varying task conditions. In general, it seems risky to assume that a single unitary process underlies conceptual processing across task conditions.

4.4. *Towards a theory of property generation*

As we have seen, responses appear to originate in two systems during property generation to word cues. Initially, properties originate in word association, but then increasingly originate in situated simulation. Clearly, well-specified accounts of these two processes and their interaction remain to be developed.

Elsewhere, we describe in greater detail how the simulation process could produce verbal responses that describe the properties of a concept (Barsalou 2003a, 2005a; Wu and Barsalou 2009). To summarize this account briefly, we assume that simulators represent components of experience in much the same way that traditional concepts represent categories (Barsalou 1999). We further assume that a simulator is typically associated with a word. The word “beak,” for example, is associated with simulator for *beak*, integrating the multi-modal information associated with this property. Once property simulators and associated words develop for a category, they can construe a simulated category instance, thereby generating properties. On being asked to generate *bird* properties, for example, a person might first simulate a particular *bird* exemplar (e.g. a particular *owl*). Once this simulation is active, the person might then use property simulators to construe its space-time regions, such as regions that contain *beak*, *feathers*, *fly*, *hoot*, and so forth. As each property simulator is bound to its corresponding region, the word associated with the simulator be-

comes active and may then be produced as a behavioral response. Over the course of interpreting the *owl* simulation in this manner, a set of properties is generated verbally.

Again, however, we assume that property names can be produced simply as word associates of the cue, bypassing the process of interpreting a simulation (e.g. “bird” activates “beak” directly). Nevertheless, we assume that properties are typically generated in both manners, as we have argued here. We further assume that these processes interact, such that generating a word in turn generates a corresponding simulation, and conversely, that generating a simulation in turn produces associated words. An important goal for future research is to specify this process in greater theoretical detail, and to develop analytic experiments that assess and develop such accounts further.

Appendix A: Response coding scheme

The response categories below are ordered from most to least linguistic (and, conversely, from least to most conceptual). When two or more categories were possible for coding a response, the most linguistic one was used. The general category of linguist responses included the specific categories of CF, BF, SS, RS, SN, and AN below. The general category of taxonomic responses included the specific categories of DH, DL, and DS. The general category of object-situation responses was identical to the specific category of OS. See text for further description.

CF Compound continuation forward

These word associates produce larger compound phrases that are common in English and that may be lexicalized. A forward continuation is one where the association follows after the cue in the continuation. Each of these must be a common phrase to be included as a continuation, not just something that it is possible to say in English. The participant may just say part of the continuation (e.g. BEE → sting), or say the whole compound (e.g. BEE → bee sting).

Examples: BEE → sting, hive; GOLF → club, course

CB Compound continuation backward

These word associates produce larger compound phrases that are common in English and that may be lexicalized. A backward continuation is one where the association precedes the cue in the continuation. Each of these must be a common phrase to be included as a continuation, not just something that it is possible to say in English. The participant may just say part of the continuation (e.g. BEE → honey), or say the whole compound (e.g. BEE → honey bee).

Examples: BEE → honey; GOLF → miniature

SS Sound similarity

These word associates are “clang” sound-based associations. They could be rhymes or also words that sound very similar.

Examples: BUMPY → lumpy; ACCEPT → except

RS Root similarity

These word associates have the same root morpheme as the cue, either adding or deleting morphemes. Typically, the grammatical class of the cue and the association differ (e.g. a verb will produce a noun).

Examples: ACCEPT → acceptance

SN Synonym

These word associates have essentially the same meaning as the cue. They are not close taxonomically related concepts, which were usually coded instead as DS (i.e. different concepts at the same level of a taxonomy). There may be more than one synonym for a cue, especially for actions and abstract concepts.

Examples: CAR → automobile

AN Antonym

These word associates have the opposite meaning of the cue. There may be more than one antonym for a cue.

Examples: ACCEPT → give, reject, deny, decline

DH Domain higher level category

These word associates are higher level categories in taxonomies that include the cue as a sub-category. Higher level categories could come from a variety of levels, and there could be more than one.

Examples: BEE → insect, bug, living thing

DL Domain lower level category

These word associates are lower level categories in taxonomies that are sub-categories of the cue. Lower level categories could come from a variety of levels, including specific individuals, and there could be more than one.

Examples: CAR → sedan; DOLL → Raggedy Ann; EUROPE → France

DS Domain same level category

These word associates are contrasting categories at the same level of a taxonomy or semantic field, having common superordinate or domain, although this may be difficult to discern for actions and abstract concepts. There could be many contrasting concepts at the same taxonomic level for a cue, which may be fairly diverse and vaguely related for actions and events.

Examples: BEE → wasp, ant, grasshopper; ACCEPT → take, acknowledge, forgive

OS Object or Situation Descriptor

These word associates are a property of the cue concept or are some other associated aspect of the same situation, including settings, thematic objects, actions, events, goals, mental states, etc.

Examples: BEE → wings, summer, flowers; GOLF → sunshine, boring, Jack Nicklaus

N None

The responses did not fit into any of the other ten coding categories.

Appendix B: Cue words in Experiments 1 and 2**Forward continuation**

taxi, golf, golden, federal, financial, fashion, self*, extension*

Backward continuation

muffin, league, cane, station, doll, lamp, lens, disc

Synonym

teacher, guess, calculate, breakable, throw, thief, secluded, excellent*

Antonym

heavy, guilty, past, exit, accept, bumpy, divorce, good*

Taxonomic category

car, bee, cello, tiger, lobster, cashew, chimpanzee, parsley

Semantic field

winter, September, north, Jupiter, Monday, monthly, gallon, Europe

Stereotypical object property

wings, stubborn, smelly, slow, horns, slimy, sweet, green

Brand

Crest, Tylenol, Nike, Budweiser, Speedstick, Marlboro, Microsoft, McDonalds

* Not used in Experiment 2

References

- Andrews, M., G. Vigliocco & D. Vinson. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological Review* 116. 463–498.
- Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22. 577–660.
- Barsalou, L. W. 2003a. Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London: Biological Sciences* 358. 1177–1187.

- Barsalou, L. W. 2003b. Situated simulation in the human conceptual system. *Language and Cognitive Processes* 18. 513–562.
- Barsalou, L. W. 2005a. Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (eds.), *Building object categories* (Carnegie Symposium Series), 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*, 9–42. New York: Cambridge University Press.
- Barsalou, L. W. 2008c. Situating concepts. In P. Robbins & M. Aydede (eds.), *Cambridge handbook of situated cognition*, 236–263. New York: Cambridge University Press.
- Barsalou, L. W., C. Breazeal & L. B. Smith. 2007. Cognition as coordinated non-cognition. *Cognitive Processing* 8. 79–91.
- Barsalou, L. W., P. M. Niedenthal, A. Barbey & J. Ruppert. 2003. Social embodiment. In B. Ross (ed.), *The psychology of learning and motivation*, Volume 43, 43–92. San Diego: Academic Press.
- Barsalou, L. W., A. Santos, W. K. Simmons & C. D. Wilson. 2008. Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg & A. C. Graesser (eds.), *Symbols, embodiment, and meaning*, 245–283. Oxford: Oxford University Press.
- Barsalou, L. W. & K. Wiemer-Hastings. 2005. Situating abstract concepts. In D. Pecher & R. Zwaan (eds.), *Grounding cognition: The role of perception and action in memory, language, and thought*, 129–163. New York: Cambridge University Press.
- Burgess, C. & K. Lund. 1997. Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes* 12. 177–210.
- Chaffin, R. 1997. Associations to unfamiliar words: Learning the meanings of new words. *Memory and Cognition* 25. 203–226.
- Collins, A. M. & E. F. Loftus. 1975. A spreading activation theory of semantic processing. *Psychological Review* 82. 407–428.
- Craik, F. I. M. 2002. Levels of processing: Past, present . . . and future? *Memory* 10. 305–318.
- Craik, F. I. M. & R. S. Lockhart. 1972. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior* 11. 671–684.
- Craik, F. I. M. & E. Tulving. 1975. Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General* 104. 268–294.
- Cree, G. S. & K. McRae. 2003. Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General* 132. 163–201.
- Damasio, A. R. 1989. Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33. 25–62.
- Fodor, J. A. 1975. *The language of thought*. Cambridge, MA: Harvard University Press.
- Gigerenzer, G. 2000. *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.
- Glaser, W. R. 1992. Picture naming. *Cognition* 42. 61–105.
- Glenberg, A. M. 1997. What memory is for. *Behavioral and Brain Sciences* 20. 1–55.
- Hampton, J. A. 1979. Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior* 18. 441–461.
- James, C. 1975. The role of semantic information in lexical decisions. *Journal of Experimental Psychology: Human Perception & Performance* 104. 130–136.

- Jordens, S. & S. Becker. 1997. The long and short of semantic priming effects in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23. 1083–1105.
- Kan, I. P., L. W. Barsalou, K. O. Solomon, J. K. Minor & S. L. Thompson-Schill. 2003. Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge. *Cognitive Neuropsychology* 20. 525–540.
- Kosslyn, S. M. 1976. Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval times. *Memory and Cognition* 4. 291–297.
- Landauer, T. K. & S. T. Dumais. 1997. A solution to Plato's Problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review* 104. 211–240.
- Lockhart, R. S. 2002. Levels of processing, transfer-appropriate processing, and the concept of robust encoding. *Memory* 10. 397–403.
- Louwerse, M. M. 2008. Embodied representations are encoded in language. *Psychonomic Bulletin and Review* 15. 838–844.
- Martin, A. 2001. Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone (eds.), *Handbook of functional neuroimaging of cognition*, 153–186. Cambridge, MA: MIT Press.
- Martin, A. 2007. The representation of object concepts in the brain. *Annual Review of Psychology* 58. 25–45.
- McRae, K., G. S. Cree, M. Seidenberg & C. McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods, Instruments, and Computers* 37. 547–559.
- Morris, C. D., J. D. Bransford & J. J. Franks. 1977. Levels of processing versus test-appropriate strategies. *Journal of Verbal Learning and Verbal Behavior* 16. 519–533.
- Murphy, G. L. 2002. *The big book of concepts*. Cambridge, MA: MIT Press.
- Nelson, D. L., C. L. McEvoy & T. A. Schreiber. 1999. The University of South Florida word association, rhyme and word fragment norms. <http://web.usf.edu/FreeAssociation/>
- Paivio, A. 1971. *Imagery and verbal processes*. New York: Holt, Rinehart, & Winston.
- Paivio, A. 1986. *Mental representations: A dual coding approach*. New York: Oxford University Press.
- Pylyshyn, Z. W. 1984. *Computation and cognition*. Cambridge, MA: MIT Press.
- Rosch, E. & C. B. Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7. 573–605.
- Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson & P. Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive Psychology* 8. 382–439.
- Shulman, H. G. & T. C. B. Davidson. 1977. Control properties of semantic coding in the lexical decision task. *Journal of Verbal Learning and Verbal Behavior* 16. 91–98.
- Simmons, W. K., S. B. Hamann, C. N. Harenski, X. P. Hu & L. W. Barsalou. 2008. fMRI evidence for word association and situated simulation in conceptual processing. *Journal of Physiology—Paris* 102. 106–119.
- Solomon, K. O. & L. W. Barsalou. 2004. Perceptual simulation in property verification. *Memory and Cognition* 32. 244–259.
- Stone, G. O. & G. C. Van Orden. 1993. Strategic control of processing in word recognition. *Journal of Experimental Psychology: Human Perception and Performance* 19. 744–774.
- Smith, E. E. 1978. Theories of semantic memory. In W. K. Estes (ed.), *Handbook of learning and cognitive processes*, Volume 6, 1–56. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Smith, E. E. & D. L. Medin. 1981. *Categories and concepts*. Cambridge, MA: Harvard University Press.

- Thompson-Schill, S. L. 2003. Neuroimaging studies of semantic memory: inferring “how” from “where”. *Neuropsychologia* 41. 280–292.
- Tulving, E. & D. M. Thomson. 1973. Encoding specificity and retrieval processes in episodic memory. *Psychological Review* 80. 352–373.
- Wilson-Mendenhall, C. D., L. F. Barrett, W. K. Simmons & L. W. Barsalou. 2011. Grounding emotion in situated conceptualization. *Neuropsychologia*.
- Wilson-Mendenhall, W. K., Simmons, A. Martin & L. W. Barsalou. Under review. Contextual processing of abstract concepts reveals neural representations of relevant non-linguistic content.
- Winer, B. J. 1971. *Statistical principles in experimental design*, 2nd edn. New York: McGraw-Hill.
- Wu, L. L. & L. W. Barsalou. 2009. Perceptual simulation in conceptual combination: Evidence from property generation. *Acta Psychologica* 132. 173–189.
- Yap, M. J., D. A. Balota, M. J. Cortese & J. M. Watson. 2006. Single versus dual process models of lexical decision performance: Insights from RT distributional analysis. *Journal of Experimental Psychology: Human Perception and Performance* 32. 1324–1344.
- Yeh, W. & L. W. Barsalou. 2006. The situated nature of concepts. *American Journal of Psychology* 119. 349–384.