Activating event knowledge

Mary Hare [a,*], Michael Jones [b], Caroline Thomson [c], Sarah Kelly [c], Ken McRae [c]

[a] Department of Psychology, Bowling Green State University, Bowling Green, OH 43403, USA
[b] Indiana University, Bloomington, IN, USA
[c] University of Western Ontario, London, Ontario, Canada

ARTICLE INFO

ABSTRACT

An increasing number of results in sentence and discourse processing demonstrate that comprehension relies on rich pragmatic knowledge about real-world events, and that incoming words incrementally activate such knowledge. If so, then even outside of any larger context, nouns should activate knowledge of the generalized events that they denote or typically play a role in. We used short stimulus onset asynchrony priming to demonstrate that (1) event nouns prime people (*sale–shopper*) and objects (*trip–luggage*) commonly found at those events; (2) location nouns prime people/animals (*hospital–doctor*) and objects (*barn–hay*) commonly found at those locations; and (3) instrument nouns prime things on which those instruments are commonly used (*key–door*), but not the types of people who tend to use them (*hose–gardener*). The priming effects are not due to normative word association. On our account, facilitation results from event knowledge relating primes and targets. This has much in common with computational models like LSA or BEAGLE in which one word primes another if they frequently occur in similar contexts. LSA predicts priming for all six experiments, whereas BEAGLE correctly predicted that priming should not occur for the instrument–people relation but should occur for the other five. We conclude that event-based relations are encoded in semantic memory and computed as part of word meaning, and have a strong influence on language comprehension.

What information do comprehenders use as they read and understand words and sentences? As a variety of results in sentence and discourse processing demonstrate, a crucial part of this information is knowledge of common events or situations in the world (Altmann, 1999; Altmann & Kamide, 1999; Camblin, Gordon, & Swaab, 2007; Hess, Foss, & Carroll, 1995; MacDonald, 1994; McKoon & Ratcliff, 2005; Vu, Kellas, Petersen, & Metcalf, 2003). But although such information is known to be important at the sentence level, it is rarely addressed at the level of individual words. Priming studies investigating word meaning tend to focus instead on semantic relatedness, often narrowly defined as the relationship between members of the same category, such as *horse* and *cow* (Fischler, 1977; Lupker, 1984; Shelton & Martin, 1992), or on a broader set of associative rela-

tions, generally determined through a normative word association task. With the notable exception of Moss, Ostrin, Tyler, and Marslen-Wilson (1995), very few investigations into the organization of semantic memory have addressed the role of event- or situation-based relations.

This appears to be a crucial gap in the literature, because in order to understand the influence of event knowledge on comprehension, we must also understand what information is made available when specific words (or classes of words) are encountered. One fruitful method for investigating this issue is semantic priming. In this article, we present a set of priming studies designed to test whether single words activate salient aspects of the representations of real-world events. We then simulate the experiments using two co-occurrence based models, Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) and BEAGLE (Jones & Mewhort, 2007), to better understand the cause of the effects. First, however, we briefly review

* Corresponding author.
E-mail address: mlhare@bgsu.edu (M. Hare).

evidence that pragmatic knowledge of this sort plays a central role in sentence and discourse comprehension, and then summarize recent work on the role of event-based information in priming.

1. Event knowledge in comprehension

Research in sentence processing has emphasized the importance to comprehension of the thematic roles established by the verb, that is, the agents, patients, instruments, and, in some cases, locations that represent the "modes of participation" or roles that the arguments of the verb play in the event that the verb denotes (Tanenhaus & Carlson, 1989). Comprehenders rapidly compute the plausibility of a role filler in a given role, and use this information in interpreting further text. In one study on comprehension of structurally ambiguous reduced relative clauses, McRae, Spivey-Knowlton, and Tanenhaus (1998) found that sentences whose initial noun phrase was a good *patient* of the verb (e.g. *the candidate interviewed*), were relatively easy to comprehend when they continued as reduced relative clauses (*...by the reporter...*), compared to cases when that noun phrase was a good *agent*, such as *The reporter interviewed*. Similar effects have been found by Stowe (1989), among others.

Thematic role information is often described as an aspect of a verb's argument structure, but its role in comprehension goes beyond strictly linguistic knowledge to reflect the comprehender's understanding of how situations plausibly occur in the world (McKoon & Ratcliff, 2005; Tanenhaus, Carlson, & Trueswell, 1989). Boland, Tanenhaus, Garnsey, and Carlson (1995), for example, compared reading times to sentence pairs that differed only in the recipient (e.g. *which preschool nursery/military base did Hank deliver machine guns to last week?*). There was an effect of plausibility at *machine guns*, with slower reading times following *preschool nursery* than following *military base*, showing that plausibility is computed relative to the entire event being described, not simply the verb *delivered*. Altmann (1999) made this point even more strongly by taking the potential goal argument out of the sentence and presenting it earlier in a short discourse context (*Hank parked his van outside the preschool nursery. He delivered some machine guns to the military base next door.*). Although the actual recipient, *military base*, eventually turned out to be plausible, there was again a plausibility effect at *machine guns*, showing that readers anticipated a patient that would be appropriately delivered to a nursery.

These results, and many similar ones, show that comprehension relies on rich pragmatic knowledge about real-world events. Incoming words incrementally activate that knowledge, serving as cues that add to and modulate the developing representation. But if individual nouns and verbs play this role in discourse, then even outside of any larger context they should activate schematic knowledge of the generalized events that they denote or play a role in. And indeed, techniques used to test semantic relatedness among words, particularly short stimulus asynchrony (SOA) priming, have been shown to tap into this information.

In one set of studies using this paradigm, Ferretti, McRae, and Hatherell (2001) tested whether information about plausible thematic role fillers would be activated by the verb in isolation. They found that it was: Verbs primed nouns referring to good fillers of their thematic roles, including prototypical agents (*arresting* primes *cop*), patients (*serving–customer*) and instruments (*stirred–spoon*). Facilitation is also found in the opposite direction, from nouns to verbs (McRae, Hare, Elman, & Ferretti, 2005). In this set of studies, verbs were named aloud following nouns referring to agents, patients, instruments, and locations typical of the event that the verb labeled. Robust priming was found for all four relations, in both short and long SOA priming tasks.

Working from a rather different theoretical approach, Moss et al. (1995) found evidence for event-based relations as well. They conducted a number of priming studies investigating the types of relationship automatically activated upon hearing single words, two of which, *instrument* and *script* relations, are directly relevant to the issues we are addressing. Moss et al. argued that when someone reads or hears an instrument prime, they activate functional information regarding how it is used, facilitating recognition of a target referring to a typical patient of the action that the instrument performs. Script primes, which referred to a mixture of events and locations, were intended to activate general event-based knowledge related to those terms, and were paired with target nouns referring to entities commonly found at that event or location. Script priming was relatively weak in some tasks, with marginal effects for unassociated pairs in auditory single-word continuous lexical decision, and with no priming with single-word visual presentation regardless of the degree of association. However, both instrument and script relations led to priming in a paired auditory lexical decision task with a short (200 ms) inter-stimulus interval (ISI) between prime and target.

Taken together, these studies indicate that single words, whether a verb denoting a common event or situation, or a noun denoting a typical participant, suffice to rapidly activate related event knowledge. The goal of the present work is to test this in more detail. First, we examine whether nouns referring to common events, like *sale* or *accident*, prime salient participants in those events just as event verbs do. Following that, we report two sets of experiments testing whether nouns that are salient cues to classes of events or situations activate that knowledge and consequently prime other salient participants. We also argue that our results are due to higher-order semantic representations, not simple co-occurrence, and test this by simulating the experiments using LSA and BEAGLE. Finally, directionally-sensitive measures from BEAGLE suggest that, although we do not obtain priming from types of instruments to the types of people who use them, priming should be obtained in the other direction. This is shown to be the case in the final experiment.

2. Experiment 1

Experiment 1 used short SOA priming to test whether event nouns activate detailed knowledge of the event to which they refer. The primes were nouns denoting generalized events (*accident, trip*), and the targets were nouns

denoting typical people or objects involved in those events (*policeman, luggage*). In Experiment 1a, all targets referred to types of people, whereas in Experiment 1b, they were names of objects. Short SOA priming was used because it is assumed to provide a window into the organization of semantic memory, with effects relatively uncontaminated by strategic processing (de Groot, 1984; den Heyer, Briand, & Dannenbring, 1983; Neely, 1977). Furthermore, shorter SOAs better match single-word reading times in normal sentence comprehension than do longer SOAs.

Semantic decisions were used for two reasons. First, they force participants to rely on semantic information. Second, they allow us to use filler pairs that are related but require a *no* response on the decision task. This balances the related test items, so that relatedness is not a cue to a "yes" or "no" response, thus discouraging retrospective processing (McRae & Boisvert, 1998; Neely, Keefe, & Ross, 1989).

The procedures used in our studies – low relatedness proportion, semantic decision task, and short SOA – are assumed to support automatic priming effects, as noted above. We also use paired rather than single-word presentation of the prime and target: That is, rather than requiring a decision at each word, prime–target pairs are presented in sequence and a semantic decision is made only to the target. Thus prime and target are overtly paired together. We note that earlier work has raised concerns that this presentation, at least in conjunction with a lexical decision task, may encourage strategic processing in a way that single-word presentation, in which a response is required at each word, does not (Moss et al., 1995). However, McRae and Boisvert (1998) compared priming effects crossing the two presentation procedures and two tasks, lexical and semantic decision, and found no evidence that paired presentation increased strategic effects: Priming effects were in fact larger with single-word than with paired presentation. Overall, then, we are confident that the procedures used here are likely to reduce strategic effects and result in automatic priming.

In Experiment 1a, in which the test targets referred to people, we used an animacy decision task ("Does the word refer to something that is alive?"), whereas in Experiment 1b, in which the test targets were inanimate objects, the task was concreteness decision ("Does the word refer to something that is a concrete object, that is, something you can touch?"). We predicted shorter decision latencies when targets were preceded by nouns referring to events in which the person or object commonly participates, compared to when targets were preceded by unrelated event nouns.

2.1. Experiment 1a

2.1.1. Method
2.1.1.1. Participants. Twenty University of Western Ontario undergraduates participated in the priming study for course credit, and a separate 20 University of Western Ontario undergraduates participated in the norming study. In all experiments reported in this article, all participants were native speakers of English and had normal or corrected-to-normal visual acuity. No participant took part

in more than one experiment, whether priming or norming.

2.1.1.2. Materials. In all of the priming experiments, materials were selected based on production norms designed to assess comprehenders' knowledge of the types of people, animals, or objects that are involved in common events. In Experiment 1a, participants were given event nouns such as *sale* and asked to list the types of people that typically are found at each event. Space was provided for up to five responses for each item. No time limit was imposed. There were 52 items per list.

Responses in the production norms were given a weighted score based on their rank order and frequency (that is, the number of participants who listed it first through fifth). Each response's weighted score was calculated by multiplying the number of times it was produced first by 5, the number of times it was produced second by 4, and so on, and then summing these products. Items were chosen for the priming experiment based on these weighted scores. In the majority of cases, the items with the highest weighted scores were used, but there were certain constraints that occasionally eliminated the strongest response. In a few cases, the response and the prime formed a common phrase (e.g. *birthday–present*, which was a highly weighted item in the norms of Experiment 1b), and these were excluded to avoid phrasal priming. Multi-word responses were also eliminated because the priming task required a single-word target. In a few other cases, the same response had the highest score for more than one prime, but we used each target only once. In these cases, we substituted another highly-ranked response, substituted a near synonym of the best response, or did not use that prime. Following these guidelines, for Experiment 1a, we chose 18 event nouns paired with targets denoting types of people, such as *sale–shopper* or *accident–policeman* (the items are presented in Appendix A). The mean weighted score for these event–people pairs was 60 (range = 31–89; maximum of 100).

We also collected free association measures for our items. Because the majority of the stimuli appear in Nelson, McEvoy, and Schreiber's (1998) norms, the values for those items were taken from those norms (both forward associations from primes to targets, and backward association from targets to primes). For the remainder, we collected free association norms using Nelson et al.'s procedure and 40 participants per word.

Note that differences are expected between free association norms and our more constrained production norms. Responses in the free association task are influenced by a variety of factors, including knowledge of various semantic relations (multiple types of featural relations, taxonomic relations, and so on), common phrases, and phonological relations. Some of the relations that influence responses in free association norms are indeed based on people's knowledge of common events. Thus, in some ways, our production norms can be considered to be a version of a constrained free association task – we provided participants with a word stimulus and asked them to produce a linguistic response. The crucial difference, however, is that our norms targeted specific semantic relations. In Experi-

ment 1a, for example, we targeted participants' knowledge of the types of people that often participate in certain types of events. Responses thus reflected these specific relationships, rather than undifferentiated associations.

We also note that throughout this article we distinguish between semantic relations and responses in the normative association task. In essence, all of our items are associated, in the general sense of association. That is, *shopper* is associated with *sale* because shoppers tend to be found at sales, regardless of whether or not *shopper* is a strong response to *sale* in a free association task (its forward association is only .03, and its backward association is 0). Therefore, we use free association norms in two ways in this article. First, we report the free association values for the stimuli in each experiment. Second, to address views in which such free associations are primary, we conduct additional analyses for each experiment using only the items that are associated either weakly or not at all. However, while we are careful to show that normative association does not explain the experimental results, we do not draw a line between "semantics" and "association" in the more general sense. And indeed, the relationships we test here could well be structured in semantic memory as a spreading activation network (Rumelhart, Smolensky, McClelland, & Hinton, 1986). We return to this point in the General discussion.

For Experiment 1a, the mean forward association strength (that is, association from prime to target) was quite low at .06, with a minimum of 0 (when the type of person was never produced as a response to the event noun) and a maximum of .24. For one item (*appointment–doctor*), the target was a primary associate of the prime, while for 4 items, the target was never produced given the prime. Backward associative strength (i.e., from target to prime) was also quite low (*M* = .02, range = 0–.20).

2.1.1.3. Lists. Two lists were created, with each list containing half of the targets paired with related primes, and the other half paired with unrelated primes. In all experiments, unrelated prime–target pairs were created by re-pairing the related trials in the opposite list (e.g. *shopper* was preceded by *sale* in list 1 and by *accident* in list 2). The participant's task was to decide if the target word referred to a living thing, and filler trials were designed to obscure any correlation between a "yes" response and prime–target relatedness. Thus, in addition to the 9 related (*sale–shopper*) and 9 unrelated event–people pairs *(election–kids)*, there were 54 filler pairs: 9 inanimate targets preceded by related event primes (*blizzard–snow*), 9 inanimate targets preceded by unrelated event primes (*migration–piano*), 18 inanimate targets preceded by unrelated inanimate primes (*backpack–towel*) and 18 animate targets preceded by unrelated inanimate primes (*taxi–bird*). The relatedness proportion was .25. For each list, 50% of the prime–target pairs were "yes" trials and 50% of the pairs were "no" trials. The use of equal proportions of "yes" and "no" trials, the low relatedness proportion, the types of fillers, and non-event primes served to decrease any cueing of responses. Finally, there were 24 practice trials, with the same proportion of trial types as in the main experiment. No participant saw any word twice.

2.1.1.4. Procedure. Stimuli were presented on a Macintosh computer with a 15 inch color monitor, using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). Decision latencies were measured with millisecond accuracy using a CMU button box. Words were presented in lower-case 20-point "New York" font in black, in the center of a white screen. For each trial, the participant was instructed to silently read the first word presented on the computer screen, and decide as quickly and accurately as possible whether the second word referred to a living thing.

Each trial consisted of a fixation point (∗) for 250 ms, followed by the prime (*sale*) for 200 ms, a blank screen for 50 ms, and then the target (*shopper*), which remained on the screen until the participant responded. To ensure that participants were doing the task as instructed, they were also informed that at random intervals they would be asked to write the first word of the pair they had just seen.

2.1.1.5. Design. Decision latencies and the square root of the number of errors (Myers, 1979) were analyzed by two-way analyses of variance. Only trials on which participants responded correctly were included in the decision latency analyses. The factor of interest was relatedness (related vs. unrelated), which was within both participants (F_1) and items (F_2). List was included as a between-participants dummy variable and item rotation group as a between-items dummy variable to stabilize variance that may have resulted from rotating participants and items over the two lists (Pollatsek & Well, 1995).

2.1.2. Results

Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (1% of trials). Mean decision latency and percent errors are presented for each condition in Table 1. Decision latencies were 32 ms shorter when the target noun was preceded by a related than by an unrelated event noun, $F_1(1,18) = 5.30$, $p < .04$, $\eta^2 = .23$; $F_2(1,16) = 7.74$, $p < .02$, $\eta^2 = .33$. There were no reliable differences in error rates, $F_1(1,18) = 2.00$, $p > .1$, $F_2 < 1$.

Although normative word association was quite low for these items, we conducted further analyses on decision latencies to determine whether priming was obtained for items that are associated weakly or not at all. For these

Table 1
Mean decision latencies (ms) and percent errors for Experiment 1. Event nouns.

	Event–people		Event–things	
	M	*SE*	*M*	*SE*
Decision latency				
Unrelated	622	18	771	47
Related	590	17	738	43
Facilitation	32[*]		33[*]	
Percent errors				
Unrelated	2.2	1.0	2.7	1.2
Related	1.1	0.8	2.3	1.0
Facilitation	1.1		0.4	

[*] Significant by participants and items.

analyses, we chose only those items that were not primary associates, and had forward association strengths of less than .1. These values were chosen because it has been reported that this type of item does not lead to priming, at least when free association values alone are considered (Anaki & Henik, 2003). Fifteen prime–target pairs fit these constraints (forward association: *M* = .04, backward association: *M* = .03). Decision latencies were 40 ms shorter when the target noun was preceded by a related (593 ms) than by an unrelated event noun (633 ms), $F_1(1,18) = 7.20$, $p < .02$, $\eta^2 = .29$; $F_2(1,13) = 9.50$, $p < .009$, $\eta^2 = .42$.

2.2. Experiment 1b

2.2.1. Method
2.2.1.1. Participants. Twenty University of Western Ontario undergraduates participated in the priming study for course credit.

2.2.1.2. Materials. For the production task, 20 undergraduate students from the University of Western Ontario were given event nouns such as *picnic* and asked to list the types of objects/things that typically are found at each event. There were 52 items per list. The method and scoring for the norming task were identical to Experiment 1a. From these norms, we chose 26 events paired with targets denoting types of objects, such as *recess–ball* or *picnic–blanket*. The mean weighted score was 57 (range = 31–85; maximum of 100; items are presented in Appendix B).

The mean forward association strength was .11, with a minimum of 0 and a maximum of .63. For 4 items, the target was a primary associate of the prime, and for 5 items, it was never produced given the prime. Backward associative strength was generally low (*M* = .03, range = 0–.36).

2.2.1.3. Lists. Two lists were created, with each containing half of the targets paired with related primes, and the other half of the targets paired with unrelated primes. In addition to the 13 related event-thing pairs (*recess–ball)* and 13 unrelated event-thing pairs *(robbery–hamburger)*, there were 78 filler pairs: 13 abstract targets preceded by related event primes (*anniversary–love*), 13 abstract targets preceded by unrelated event primes (*departure–years*), 26 abstract targets preceded by unrelated concrete primes (*backpack–apathy*) and 26 concrete targets preceded by unrelated concrete primes (*elevator–banana*). As in Experiment 1a, the proportion of related prime–target pairs was .25, and 50% of the prime–target pairs were "yes" trials and 50% of the pairs were "no" trials. Finally, there were 24 practice trials, with the same proportion of trial types as in the main experiment. No participant saw any word twice.

2.2.1.4. Procedure. The procedure was identical to Experiment 1a except that participants were asked to indicate, as quickly and accurately as possible, whether or not the target word referred to a concrete thing, defined as something that can be touched, with examples like *textbook* or *milk* as opposed to abstract concepts like *equality* or *happiness*.

2.2.1.5. Design. The Design was identical to Experiment 1a.

2.2.2. Results
Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (1% of trials). Mean decision latency and percent errors are presented for each condition in Table 1. Decision latencies were 33 ms shorter when the target noun was preceded by a related event noun than by an unrelated event noun, $F_1(1,18) = 7.74$, $p < .02$, $\eta^2 = .30$; $F_2(1,24) = 4.71$, $p < .05$, $\eta^2 = .16$. There were no reliable differences in error rates, $F_1 < 1$, $F_2 < 1$.

We again conducted analyses on decision latencies to determine whether priming is obtained for items that are weakly normatively associated. There were 16 items that were not primary associates and had forward association strengths of less than .1 (forward: *M* = .02, backward: *M* = .05). Decision latencies were 46 ms shorter when the target noun was preceded by a related (723 ms) than by an unrelated event noun (769 ms), $F_1(1,18) = 19.02$, $p < .0005$, $\eta^2 = .51$; $F_2(1,14) = 6.10$, $p < .03$, $\eta^2 = .30$.

2.2.3. Discussion
Experiment 1 demonstrates that event nouns prime salient aspects of the class of events that they label. Ferretti et al. (2001) found the same to be true of event verbs. But while priming from verbs to their typical role fillers might conceivably be due to activation of information in the verb's lexical representation, that cannot be the case here, for nouns are not typically assumed to assign thematic roles. Instead, these results strongly suggest that words denoting events – whether nouns or verbs – activate information about the event and its typical participants.

Both the noun priming found here and the verb priming of Ferretti et al. (2001) are expected if event memory is organized so that relatively detailed knowledge of generalized events can be computed quickly and easily from multiple cues. These results make sense when normal sentence comprehension is considered. When a noun denoting an event is read or heard, it is likely that the common participants in those events will appear downstream in the discourse, and consequently event-based expectations regarding upcoming propositions will be computed. This might also be viewed as an important aspect of constructing a situation model (Morrow, Greenspan, & Bower, 1987; Zwann & Radvansky, 1998). When an event noun is encountered, the constructed mental model will tend to involve typical participants and the roles they play in that event.

McRae et al. (2005) found that nouns denoting typical agents, patients, instruments, or locations activated well-learned event knowledge, priming verbs that overlapped substantially with the activated event space. This suggests that such typical event participants should prime other event participants as well. In Experiments 2 and 3, we investigated whether this is the case, using location and instrument nouns as primes.

3. Experiment 2

One salient piece of knowledge about many common events and situations is the location at which they com-

monly occur. Experiment 2 examines whether reading a location noun results in rapid activation of information about events that commonly occur at that location, facilitating the processing of typical event participants. Both our own previous work (McRae et al., 2005), and work by Moss et al. (1995) suggest that this should be the case. Using both auditory and visual lexical decision tasks, Moss and colleagues tested for priming for what they referred to as *script* relations. Their items included location, event, and other nouns as primes, and nouns from various semantic classes as targets. Because Moss et al. included both event and location primes, it is not entirely clear whether the priming was due to event nouns, location nouns, or a combination of the two. Furthermore, results were variable for these items: Priming was found in a paired auditory lexical decision task (Experiment 1), but when single-word presentation was used, the priming effects for non-associated script pairs were marginal with auditory presentation (Experiment 2) and nonsignificant with visual presentation, regardless of the degree of association between prime and target (Experiment 3). Nonetheless, the results are suggestive, and it is likely that any weakness in the priming effects reflected the fact that the authors relied on intuition to choose items. In our studies, we tested these relations using normatively-chosen item pairs and the same priming paradigm as in Experiment 1. In Experiment 2, the prime was a noun denoting a location, and the target a noun denoting either a type of person or animal (Experiment 2a), or a type of object (Experiment 2b) commonly found at that location.

3.1. Experiment 2a

3.1.1. Method
3.1.1.1. Participants. Twenty-four University of Western Ontario undergraduates participated in the priming study for course credit. Two participants were dropped due to exceptionally slow decision latencies, leaving 11 participants per list.

3.1.1.2. Materials. Target stimuli were again chosen on the basis of production norms. All aspects of the norming except for the materials were the same as in Experiment 1. Twenty participants were given a list of 58 locations and were asked to list the types of people or animals typically found at each. From these norms, we chose 24 location nouns paired with animate targets (types of people or animals, e.g. *mall–shopper, barn–cow*, see Appendix C). The mean score for the location–animate pairs was 63 (range = 25–97, maximum of 100).

The mean forward association strength was .06, ranging from 0 to .43. For 2 items, the target was a primary associate of the prime, and for 11 items, it was never produced given the prime. Backward associative strength was low (M = .08, range = 0–.29).

3.1.1.3. Lists. Two lists were created containing 12 of the targets paired with related primes, and the other 12 paired with unrelated primes. The composition of the lists in terms of the proportions of various types of filler trials was the same as in Experiment 1a.

3.1.1.4. Procedure and Design. The Procedure and Design were identical to Experiment 1a, which also used an animacy decision task.

3.1.2. Results
Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (2% of trials). Mean decision latency and percent errors are presented in Table 2. Decision latencies were 37 ms shorter when targets were preceded by a related location than when they were preceded by an unrelated location, $F_1(1,20) = 4.39$, $p < .05$, $\eta^2 = .18$; $F_2(1,22) = 5.29$, $p < .04$, $\eta^2 = .19$. There were no reliable differences in error rates, both F's < 1.

There were 18 items that were not primary associates, and had forward association strengths of less than .1 (forward: M = .01, backward: M = .08). Decision latencies were 33 ms shorter when the target noun was preceded by a related (730 ms) than by an unrelated location noun (763 ms), which produced a marginal priming effect, $F_1(1,20) = 3.50$, $p < .08$, $\eta^2 = .15$; $F_2(1,16) = 2.58$, $p < .13$, $\eta^2 = .14$.

3.2. Experiment 2b

3.2.1. Method
3.2.1.1. Participants. Twenty University of Western Ontario undergraduates participated in the priming study.

3.2.1.2. Materials. For the production norms, 20 participants were given a list of 62 locations and asked to list the types of things typically found at each. From these norms, we chose 30 location nouns paired with inanimate targets (e.g. *sandbox–shovel, alley–garbage*, see Appendix D). The mean score for the location–thing pairs was 58 (range = 38–80; maximum of 100). Two lists of items were created in a manner identical to that in Experiment 1b.

The mean forward association strength was reasonably high at .17, ranging from 0 to .76. For 10 items, the target was a primary associate of the prime, and for 6 items, it was never produced given the prime. Backward associative strength was low (M = .04, range = 0–.34).

3.2.1.3. Procedure and design. The Procedure and Design were identical to Experiment 1b, in which a concreteness decision task was also used.

Table 2
Mean decision latencies (ms) and percent errors for Experiment 2: Location nouns

	Location-living		Location–things	
	M	SE	M	SE
Decision latency				
Unrelated	765	44	675	26
Related	728	36	646	28
Facilitation	37[*]		29[*]	
Percent errors				
Unrelated	3.0	1.2	3.3	1.3
Related	2.3	1.2	3.3	1.5
Facilitation	0.7		0	

[*] Significant by participants and items.

3.2.2. Results

Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (1% of trials). Mean decision latency and percent errors are presented in Table 2. Decision latencies were 29 ms shorter when the targets were preceded by a related location than when they were preceded by an unrelated location, $F_1(1,18) = 8.60$, $p < .009$, $\eta^2 = .32$; $F_2(1,28) = 5.16$, $p < .04$, $\eta^2 = .16$. There were no reliable differences in error rates, both F's < 1.

There were 16 items that were not primary associates and had forward association strengths of less than .1 (forward: $M = .03$, backward: $M = .04$). Decision latencies were 79 ms shorter when the target noun was preceded by a related (679 ms) than by an unrelated location noun (758 ms), $F_1(1,18) = 10.09$, $p < .006$, $\eta^2 = .36$; $F_2(1,14) = 41.56$, $p < .0002$, $\eta^2 = .75$.

3.2.3. Discussion

In Experiment 2, location words were used as primes, and either humans or animals (in Experiment 2a), or inanimate objects (Experiment 2b) commonly found at those locations were used as targets. Priming was obtained in both cases, suggesting that location nouns activate event knowledge during comprehension, thus activating information regarding people and objects that typically are found at that location. Locations are excellent cues to situation models (Zwann & Radvansky, 1998). In fact, Lancaster and Barsalou (1997) found that people organize short narratives in terms of multiple components of events. Although activity (that is, the verb or an event noun) is of course an important component, location, time, and central participants were shown to be important as well. Lancaster and Barsalou's results, in combination with Experiment 2 and McRae et al.'s (2005) finding that location nouns prime verbs denoting an activity that commonly takes place at those locations, demonstrate that semantic memory is organized to allow for rapid access of event-based knowledge given a location.

In Experiments 1a, 1b and 2a, the analyses of the prime–target pairs that were not normatively associated showed priming effects of a roughly equivalent magnitude to the analyses using all pairs. However, the priming effect was substantially larger in Experiment 2b for items that were not normatively associated (79 ms vs. 29 ms), which is somewhat surprising. A plausible explanation is that Experiment 2b is the one case in which the unrelated decision latencies were much longer for the items that were not normatively associated (758 ms vs. 675 ms). Priming is a facilitation effect, and therefore the more difficult the targets, the greater the opportunity for observing facilitation. As such, we do not place any theoretical significance on this difference in priming effects.

4. Experiment 3

Instruments are a salient aspect of many events. McRae et al. (2005) found that instrument nouns prime verbs denoting a class of events in which the instrument is typically used. In Experiment 3, we test whether nouns that name instruments can in turn activate aspects of those events or classes of event, therefore priming other typical event participants. We again distinguish between animate and inanimate target nouns. The instrument–people items pair an instrument with the types of people who typically use it (*wrench–plumber*). The instrument–thing items (*oven–cookies*), similar to those used in Moss et al. (1995), pair an instrument with an example of the type of item it typically acts upon. We predicted shorter decision latencies to target nouns following a related versus an unrelated prime.

4.1. Experiment 3a

4.1.1. Method

4.1.1.1. Participants. Twenty-seven undergraduates from the University of Western Ontario participated for course credit. One participant was dropped due to exceptionally slow decision latencies.

4.1.1.2. Materials. For the production norms, 20 participants were given 49 instrument nouns such as *crayon* and *hose* and asked, "What people commonly use each of the following?" From these norms, we chose 24 instrument nouns paired with targets referring to types of people (e.g. *crayon–child*, *wrench–plumber*, see Appendix E). The mean score for the instrument–people pairs was 58 (range = 38–89; maximum of 100). Two lists of items were created in an identical manner to Experiments 1a and 2a.

The mean forward association strength was .06, ranging from 0 to .29. For 1 item, the target was a primary associate of the prime, and for 11 items, it was never produced given the prime. Backward associative strength was low ($M = .03$, range = 0–.20).

4.1.1.3. Procedure and design. The Procedure and Design were identical to Experiments 1a and 2a, in which an animacy decision task was also used.

4.1.2. Results

Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (1% of trials). Mean decision latency and percent errors are presented in Table 3. Decision latencies did not differ when the target words were preceded by a related versus an unrelated instrument prime, $F_1(1,24) = 0.31$, $p > .5$, $\eta^2 = .01$; $F_2(1,22) = 0.28$, $p > .6$, $\eta^2 = .01$. There were no reliable differences in error rates, both F's < 1. Note that we have replicated this null effect, using the same items, but with participants from Bowling Green State University.

There were 19 items that were not primary associates, and had forward association strengths of less than .1 (forward: $M = .01$, backward: $M = .02$). There remained no priming effect: related (768 ms), unrelated (769 ms), $F < 1$, $\eta^2 < .01$ in both analyses.

4.2. Experiment 3b

4.2.1. Method

4.2.1.1. Participants. Sixteen undergraduates from Bowling Green State University participated for course credit.

Table 3
Mean decision latencies (ms) and percent errors for Experiments 3 and 4: Instrument nouns.

	Instrument–people		Instrument–things		People–instruments	
	M	SE	M	SE	M	SE
Decision latency						
Unrelated	756	33	793	44	776	28
Related	766	37	735	37	731	26
Facilitation	−10		58[*]			45[*]
Percent errors						
Unrelated	3.2	0.9	1.0	0.7	3.6	1.3
Related	3.5	1.1	2.1	1.2	2.6	1.1
Facilitation	−0.3		−1.1		1.0	

[*] Significant by participants and items.

4.2.1.2. Materials. For the production norms, 20 participants were given a list of 48 instrument nouns and asked "What things do people commonly act upon with each of the following?" From these norms, we chose 24 instruments paired with inanimate targets (e.g. *razor–face, oven–cookies*, see Appendix F). The mean weighted score for the instrument–object pairs was 64 (range = 41–94; maximum of 100). Two lists of items were created in an identical manner to Experiment 1b.

The mean forward association strength was quite high at .13, ranging from 0 to .63. For 5 items, the target was a primary associate of the prime, and for 5 others, it was never produced given the prime. Backward associative strength was low (*M* = .01, range = 0–.19).

4.2.1.3. Procedure and design. The procedure and design were identical to Experiments 1b and 2b, in which a concreteness decision task was also used.

4.2.2. Results

Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (1% of trials). Mean decision latency and percent errors are presented in Table 3. Decision latencies were 58 ms shorter when the target words were preceded by a related versus unrelated instrument prime, $F_1(1,14) = 9.59$, $p < .008$, $\eta^2 = .41$; $F_2(1,22) = 10.72$, $p < .004$, $\eta^2 = .33$. There were no reliable differences in error rates, both *F*'s < 1.

There were 12 items that were not primary associates and had forward association strengths of less than .1 (forward: *M* = .02, backward: *M* = .003). Decision latencies were 73 ms shorter when the target noun was preceded by a related (729 ms) than by an unrelated instrument noun (802 ms), $F_1(1,14) = 8.58$, $p < .02$, $\eta^2 = .38$; $F_2(1,10) = 8.92$, $p < .02$, $\eta^2 = .47$.

4.2.3. Discussion

In Experiment 3, instruments primed objects on which they are typically used, but not types of people who typically use them. Given that both relationships should be available to the comprehender, the discrepancy in the results stands to be explained.

Simple methodological factors can be ruled out. For example, because priming effects are generally facilitative, priming may not be obtained when target words are short

or high in frequency, two properties that facilitate processing as well. However, if this had been the case with our items, it should have been reflected in the unrelated decision latencies, and it was not. The mean unrelated decision latency for the instrument–people pairs (756 ms) is easily within the range of the other conditions in which priming was found. Event–people (622 ms) and location–things (675 ms) have substantially shorter unrelated decision latencies. The means for location–animates (765 ms) and event–things (771 ms) are quite similar to location–people, and the instrument–things mean unrelated decision latency is somewhat longer (793 ms).

Error rates for instrument–people were in the range of the other conditions as well (3.2% for instrument people; range of other conditions was 1.0–3.3%). Nor can the answer be found in the strength of the production norms, as the instrument–people condition was similar to the others in which priming was obtained. Furthermore, although association strengths were quite low for the instrument–people pairs (forward = .06, backward = .03), they were essentially the same as event–people (.06 and .02) and location–people (.06 and .08), both of which showed priming. Also note that the 26 participants used in Experiment 3a (which was a greater number than in any of the other experiments) gives power equal to .88 to detect an effect that is the size of the average of those found in the other experiments. However, if the effect size was only half of those found in the other experiments, 80 participants would be required to obtain power equal to .8. In addition, the instrument–people priming effect was actually 10 ms in the wrong direction, and we have replicated this null effect. For all of these reasons, we believe that the lack of an instrument–people priming effect is unlikely to be a Type II error.

Instead, we propose that the instrument–person relationship, unlike the others, was not sufficiently constrained to support short SOA priming. Many of the common instruments used in Experiment 3, like common instruments in general, are used by many types of people. Although a knife is indeed used by a chef, for example, it is not exclusive to chefs; many others commonly use knives as well. Given the generality of instruments like *rag, pen,* and *hose* across agents, the specific relationships we tested may not have been sufficiently constrained to produce short SOA priming effects. As a post hoc test of this explanation, we selected a subset of seven instruments that, based on our intuition,

were highly specific to types of individuals who might use them (*microscope*, *saddle*, *rod*, *compass*, *needle*, *stopwatch*, and *ladder*). When only these seven items were analyzed, there was a 64 ms priming effect (related = 758 ms; unrelated = 822 ms). Due to the small number of items, this effect was not significant, $F_2(1,5) = 3.02$, $p > .1$. However, the contrast of the 64 ms priming effect for these items with the overall effect of −10 ms is certainly suggestive.

In the next section, we present corpus-based computational simulations of our priming experiments. Given the similarity between the models used and our own account, measures from the simulations can offer an independent corpus-based validation of our results. They also offer insight into the reasons we failed to find a priming effect from instruments to types of people who use them, and provide a prediction for Experiment 4.

5. Computational analyses

On our account, the observed facilitation from prime to target is a product of shared event knowledge, which is part of each word's semantic representation. This account has much in common with computational models like LSA, in which priming between two words is based on their frequency of occurrence in similar contexts. As a consequence of participating in the same event, or, in the case of event noun primes, participating in the event named by the prime, our prime–target pairs are quite likely to occur in the same contexts in text or speech. However, we suggest that both co-occurrence and priming effects arise from underlying semantic relationships, in this case knowledge of events in the world. Co-occurrence models start from established corpora and so are agnostic on the question of why particular patterns of co-occurrence arise in the first place. On the account proposed here, one important reason is that language users speak or write about their knowledge and experience of common events, and words referring to participants in the same class of events occur in similar contexts as a consequence. These co-occurrence patterns reflect the underlying meaning being expressed, and it is the semantic relationships, not the co-occurrence in itself, that drives the priming.

LSA observes the frequency with which words co-occur in documents across a text corpus, and then uses singular-value decomposition (akin to factor analysis) to identify the latent dimensions that account for the maximum variance in the frequency distribution of words over documents. A word's semantic representation in LSA is a vector pattern over these latent semantic dimensions. LSA may not be an appropriate model to test our account because it learns co-occurrences within full documents rather than sentences, which causes the model to overemphasize associative relations in priming (Jones, Kintsch, & Mewhort, 2006). In addition to LSA, we therefore used Jones and Mewhort's (2007) BEAGLE model to test our prime–target pairs because it learns co-occurrence structure from within sentences. We specifically used the simple "context only" version of BEAGLE: Each word is assigned an initial random vector, and a word's meaning is the sum of the vectors for all other words with which

it occurs in sentences across the training corpus. By this method, similarity between the semantic vectors for two words is increased each time they faithfully appear in the same sentential context, and is decreased each time they appear in different sentential contexts. This model is similar in spirit to recent random indexing models of semantic representation (Sahlgren, 2005).

The representations produced by co-occurrence models are often described in terms of a multidimensional semantic space in which similar words are located close together in space. The traditional metric used to compute word similarity in the resulting space is the cosine of the angle between the vectors representing two words (a normalized dot product). Table 4 shows cosines (simple word similarity) for our prime–target pairs from LSA and BEAGLE with both models trained on the Touchstone Applied Science Associates (TASA) corpus, a large-scale text corpus that approximates the text read across the lifespan of a college-level reader. Although there are differences in the algorithms used to learn the semantic space in each model, the primary difference is that LSA uses documents as its notion of context while BEAGLE uses sentences. The third column under each model gives the paired *t*-test associated with the priming effect for each experiment (related vs. unrelated). LSA cosines predict substantial facilitation in all of our experiments: related pairs are more similar than unrelated pairs for all experiments, including the instrument–people experiment for which we obtained a null effect (although this is the smallest predicted priming effect with LSA). This is consistent with previous findings that LSA cosines tend to predict priming effects that are stronger than those observed with humans (Jones et al., 2006).

When we move to a co-occurrence model such as BEAGLE, which is based on shared sentential context, predictions using cosines better match the pattern of results found across our experiments. Related prime–target pairs were significantly more similar than unrelated pairs in BEAGLE, except for instrument–people pairs, consistent with our null finding in Experiment 3a. Although BEAGLE cosine predicts the pattern of results found in our priming studies, it may be an overly simplistic measure of semantic relatedness because it only considers direct prime–target proximity and ignores the landscape of the semantic space between prime and target. In particular, Jones and Kintsch (2006) (see also Burgess & Lund, 2000) have shown that the number of intervening neighbors (NIN) between two words in semantic space may be a better predictor of priming data than the cosine is. The NIN considers the number of words that are more similar to the target than the prime is, and is therefore a measure of availability in semantic space. As such, it more directly tests our hypothesis that the relationship between instruments and people is less constrained than other relationships tested in the priming studies. As Table 4 shows, the results with NIN mirror the pattern shown by the cosine. For all experiments except instruments–people, related pairs have fewer intervening neighbors (NIN) than do unrelated pairs.

BEAGLE, unlike LSA, thus correctly predicts the pattern of effects in the six experiments. It also supports our

Table 4
Cosines and number of intervening neighbors (NIN) from the co-occurrence models.

	LSA cosine			BEAGLE cosine			BEAGLE (NIN)		
	Related	Unrelated	*t*	Related	Unrelated	*t*	Related	Unrelated	*t*
Event–people	.458	.046	6.31**	.318	.15	3.92**	1550	7361	−2.58*
Event–things	.341	.027	4.89**	.246	.08	4.17**	3589	16,606	−4.20**
Location-living	.482	.024	7.91**	.329	.12	5.22**	7475	17,134	−3.64**
Location–things	.343	.066	7.60**	.287	.15	5.01**	4199	9266	−2.12*
Instrument–people	.192	.039	3.83**	.128	.08	1.79	13,023	19,208	−1.43
Instrument–thing	.270	.001	5.62**	.203	.10	3.20**	1581	13,849	−3.96**

* $p < .05$.
** $p < .01$.

hypothesis that the instrument–people relationship was less constrained than the others, and therefore unable to support short SOA priming. Intuitively, the relationship seems to be more constrained in the opposite direction, suggesting that although *knife* does not prime *chef*, *chef* might well prime *knife*. Semantic measures from the model hint at this as well. With regard to NIN, although there are a large number of intervening neighbors in both directions for both the instrument–people pairs and the people–instrument pairs formed by reversing the items of Experiment 3a, there is an average of 400 fewer from people to instruments. Although the difference is still only marginally significant, $t(23) = 1.98$, $p = .06$, it does suggest that priming should be somewhat stronger in this direction. (Note that cosine is not a relevant measure for this comparison. Cosine based on any semantic space must predict the same priming effect in both directions because it is a bi-directional metric.)

Interestingly, this pattern is not observed from direct co-occurrences in the corpus on which the model was trained. For people–instruments and instruments–people, we estimated the conditional probability that the target word occurs in a sentence given that the prime word does. The conditional probability of instruments given people was actually slightly lower (.0031) than the conditional probability of people given instruments (.0072) in the raw corpus. Thus direct co-occurrence in the corpus predicts that priming from people to instruments should be weaker, if anything, than the priming from instruments to people. In contrast, the NIN measure in the semantic space suggests the reverse pattern: Instrument nouns should activate a broad range of people, so that the probability of any one being highly activated is unlikely. However, specific types of people, according to this measure, tend to activate fewer and more specific instruments, predicting that the probability of *knife* being highly activated when primed by *chef* should be greater than the reverse. We test these conflicting predictions in a final experiment.

6. Experiment 4

To this point, we have argued that event memory is organized in such a way that nouns can activate highly specific information about the events their referents commonly participate in, including information about other prototypical participants in those events. In Experiment 3a, however, we found that common instruments like *knife* did not prime nouns referring to people who tend to use them. Although common instruments are used by many types of people in many types of events, intuitively, the relationship appears to be more constrained in the opposite direction, and semantic measures from the corpus analysis suggest this as well. In a final priming study, we therefore tested whether the items used in Experiment 3a would show priming when the prime–target order was reversed.

6.1. Method

6.1.1. Participants

Sixteen undergraduates from the Bowling Green State University participated for course credit.

6.1.2. Materials

The critical items were identical to those used in Experiment 3a, except that the prime–target order of each noun pair was reversed, so that the instrument primes of Experiment 3a served as the targets in the current study. The items are listed in Appendix E. Because the target words were all inanimate, the concreteness decision task was used, and filler items were constructed as in Experiments 1b, 2b and 3b.

The mean forward association strength was low at .03, ranging from 0 to .20. For 1 item, the target was a primary associate of the prime, and for 11 items, it was never produced given the prime. Backward associative strength was low ($M = .06$, range = 0–.29).

6.1.3. Procedure and design

The procedure and design were identical to Experiment 3b, in which a concreteness decision task was also used.

6.2. Results and discussion

Decision latencies greater than three standard deviations above or below the grand mean were replaced by that value (<1% of trials). Mean decision latency and percent errors are presented in the third column of Table 3. Decision latencies were 45 ms shorter when the instrument targets were preceded by a related versus an unrelated prime, $F_1(1,14) = 35.43$, $p < .001$, $\eta^2 = .72$; $F_2(1,22) = 7.40$, $p < .02$, $\eta^2 = .25$. There were no reliable differences in error rates, both *F*'s < 1.

There were 21 items that were not primary associates and had forward association strengths of less than .1 (forward: $M = .01$, backward: $M = .05$). Decision latencies were 37 ms shorter when the target was preceded by a related (741 ms) than by an unrelated prime (778 ms), $F_1(1,14) = 10.01$, $p < .008$, $\eta^2 = .42$; $F_2(1,19) = 4.47$, $p < .05$, $\eta^2 = .19$.

Direct co-occurrence data in the corpus predicted that priming should be stronger from instruments to people than from people to instruments, the opposite of our observed priming effects. However, the NIN measure in the semantic space predicts the observed effects. The differences in the corpus analysis are not large, but demonstrate an interesting point: The semantic space predicts the correct pattern of results even though the direct co-occurrence data on which the model was trained do not. The mental space, in other words, can be richer than the direct data from which it is learned.

7. General discussion

What information is activated when reading or hearing single words? Lexical relations are often assumed to be either categorical or associative, but the results presented here show that the range is much broader than that. Using short SOA priming, we found that nouns referring to common events primed nouns referring to typical event participants, and that nouns referring to locations primed people, animals, and objects that are typically found at those locations. Instruments, which tend to be key components of the causal structure of events, primed classes of objects on which they are typically used. Finally, although there was no priming from instruments to types of people who typically use them, nouns referring to people did prime instruments that they tend to use. Furthermore, computational analyses with LSA and BEAGLE show that measures of the semantic space also predict that pragmatic, event-based relations are encoded in semantic memory, and activated as single words are read.

We believe that this type of research is an important component of furthering the understanding of the organization of semantic memory and its relationship to language comprehension in general. That is, in most priming experiments, sets of items are constructed using responses from free association norms. Although this research has furthered the field in many ways, it does not provide information about the specific types of semantic (and event-based) relations that are encoded in semantic memory because free association responses are driven by numerous types of relations. Rather, we favor an approach in which types of relations are theoretically and empirically delineated and tested individually. A major contribution of this research, therefore, is to highlight the event-based relations that are encoded with various types of noun concepts.

7.1. Sentences and discourse

What, then, is the relationship between representations at the lexical and at the sentence or discourse level? A number of studies have tested whether discourse coherence overrides lexical relations, generally defined through normative association. Some have found effects of both factors (Hess et al., 1995; Morris, 1994, Experiment 2; Garrod & Terras, 2000). In a cross-modal primed naming study, Hess et al. crossed local context – the lexical relationship between a target word (*poem*) and an earlier word or phrase in the same sentence, such as *The English major/Computer Science major* – and global context, a short paragraph consistent or inconsistent with the target word. In this case, for example, the consistent global context was about a woman the student liked, while the inconsistent context referred to a computer program. The contexts and carrier sentences (e.g. *the English major wrote the…*) were presented in the auditory modality, followed immediately by the visual presentation of the target word, *poem*, which participants named aloud. The target was named more rapidly when the context, either global or local, was consistent, leading to main effects of global and local context with no interaction. This suggests that global context might not override local, lexical, effects.

However, more recent experiments, many of which use event-related potentials (ERPs), have found that discourse context can indeed override local effects. As one example, van Berkum and colleagues have shown that words that were equally coherent in single sentences elicit different ERP responses when these sentences form part of a larger context which makes one of the two less plausible. Thus when the sentence *Jane told her brother that he was exceptionally slow/quick* was presented in isolation, *slow* and *quick* elicited N400 components of similar size. The N400 generally indexes anomaly or implausibility, and so this indicates that the two were equally plausible. However, when the same sentence was presented in a context in which Jane discovers that her brother was faster than expected, the N400 to *slow* was significantly larger than that to *quick* (Van Berkum, Brown, Zwisterlood, Kooijman, & Hagoort, 2005; Van Berkum, Hagoort, & Brown, 1999).

Thus, although the information activated by a word can be quite different when it is presented in word pairs or in a larger context, this does not necessarily imply that the word's lexical representation is limited to qualitatively different information. Instead, the results show that processing in context depends less on traditional lexical relationships than on situational or event-based information, which has not been well studied at the lexical level.

Our results, like those of Moss et al. (1995), suggest that event-based knowledge is tapped both in and out of context. What needs to be explained, then, is why word pairs reflecting event-based knowledge do not necessarily facilitate each other in sentences. One relevant example comes from Traxler, Foss, Seely, Kaup, and Morris (2000). Although in Experiment 4 we found that in isolation, typical agents prime the instruments they use, these authors found that *axe* was not read any faster in the context *The lumberjack chopped the _____*, where the subject is a typical agent of a chopping event, than in the neutral context *The young man chopped the _____*.

Although these data may appear to argue against our account, they are in fact entirely predictable given the influence of sentential context combined with structured

event-based knowledge. An *axe* typically plays the instrument role in chopping events, and other roles are filled by, for example, typical patients (*wood, tree*) and agents (*lumberjack, axe murderer*). Previous work suggests that out of context, the prime appears to activate typical fillers of all roles. Sentence structure, however, imposes its own set of constraints on interpretation. In English there is a strong correlation between structural position and thematic roles, and comprehenders exploit this information to infer the roles that participants play. Kako (2006) showed in a series of studies that readers given noun–verb–noun strings (where the nouns were actually nonsense words) consistently assigned more agent properties to the preverbal word, and more patient properties to the post-verbal one. Work in sentence processing has also shown that comprehenders anticipate specific roles based on the structural properties of the verb (Hare, Elman, Tabaczynski, & McRae, in press; Tanenhaus et al., 1989), and so would anticipate a patient, not an instrument, immediately following *chopped*. The *lumberjack chopped the axe* example thus shows that the discourse representation includes information about roles in the event, and when the participant does not fit the structurally-defined role, or fits some other role, there is no facilitation. It also shows that the information encoded with specific lexical concepts is not limited to undifferentiated associations, but it must include relational knowledge. That is, *axe* is not simply associated with *chop*; axes are used as an instrument for accomplishing chopping.

In summary, lexical semantic knowledge does indeed provide information regarding potential upcoming concepts in the linguistic stream. When people hear about an axe, this sets up expectations for hearing about a tree or wood being chopped, or about a chopping event in general, because people know what axes are used for. Our experiments demonstrate that this type of event-based relational information can be computed quite rapidly from linguistic cues. However, grammatical structure can either facilitate or inhibit expectations for certain classes of concepts because fillers of specific roles are expected in specific sentential positions. In addition, wider discourse can be used to alter, and sometimes override, event-based expectations by altering the nature of the event being described.

7.2. Alternative accounts

7.2.1. Verb-based priming

In all of the discussion thus far, we have assumed that we have measured priming from nouns to nouns. It might be possible, however, that in some cases participants interpreted the primes as verbs. Given that we have previously demonstrated event-based priming from verbs to nouns (Ferretti et al., 2001), it is important that the primes were indeed nouns in the present experiments. To obviate any concerns, we reanalyzed all significant priming effects after removing any items for which the prime might possibly be taken as a verb and for which the verb meaning would lead to the target. That is, we removed items such as *barbecue–hamburger*, but did not exclude ones such as *trip–luggage* because the verb meaning of *trip* (to fall down

accidentally) is not related to *luggage* in any way. In all cases, priming effects remained significant (we calculated F_2 only). The results are as follows: Experiment 1a, events-people, removed *lecture*, *cruise*, and *divorce*, $F_2(1,12) = 5.47$, $p < .04$, $\eta^2 = .30$ (591 ms vs. 626 ms); Experiment 1b, events-things, removed *eclipse*, *joust*, and *barbecue*, $F_2(1,21) = 4.74$, $p < .05$, $\eta^2 = .18$ (737 ms vs. 778 ms); Experiment 2a, locations–people or animals, removed *beach*, $F_2(1,21) = 5.19$, $p < .04$, $\eta^2 = .20$ (732 ms vs. 773 ms); Experiment 2b, locations–things, removed *pool*, *closet*, *market*, and *farm*, $F_2(1,24) = 5.54$, $p < .03$, $\eta^2 = .19$ (644 ms vs. 675 ms); and Experiment 3b, instruments–things, removed *hoe*, *key*, *ladle*, *net*, *pump*, and *screen*, $F_2(1,16) = 10.54$, $p < .006$, $\eta^2 = .40$ (728 ms vs. 794 ms). We also removed any potential verb targets in Experiment 4, people–instruments. This was probably not necessary because the task was to decide whether the target referred to a concrete object, but we conducted the analysis to be conservative. Priming remained significant after removing *backpack*, *hose*, *pen*, *saddle*, and *shampoo*, $F_2(1,17) = 5.39$, $p < .04$, $\eta^2 = .24$ (738 ms vs. 788 ms). In summary, all of our items were chosen so that the noun meaning was much more salient than a potential verb meaning of the primes. After conservatively excluding any items that might have been read as verbs, priming was still obtained. Therefore, we can confidently conclude that the present experiments measured noun–noun facilitation.

7.2.1. Normative association

One issue that often arises in lexical priming studies is the role of association in activating lexical representations. This issue is often portrayed as the distinction between semantic and associative relationships, but we believe that this is a false dichotomy, one that holds only if "semantic" is defined rigidly – based, for example, on category structure – and "associated" is defined in terms of a high probability of response on a normative free association task (see Hutchinson, 2003; McNamara, 2005; and Steyvers, 2000; for similar arguments.).

We have shown that normative association does not predict our experimental results. We have also demonstrated, using a simple learning algorithm instantiated by the BEAGLE model that the mental space can be richer than the data from which it is learned. That is, statistical learning is not the same as learning statistics. Although the model learns *from* corpus co-occurrence, the resulting representation contains much more – only a fraction of what is learned comes from direct co-occurrence. As one example of this, recall that conditional probability suggests that priming should be more likely for our instrument–people than for our people–instrument pairs, which is the reverse of the human results. As another example, synonyms rarely co-occur, but the representations they develop are close in the semantic space because they occur in similar sentences. Hence the effects are not necessarily due to textual co-occurrence.

7.2.2. Mediated priming

One fundamental question in the literature on semantic memory is whether free association norms adequately measure semantic/associative structure (McKoon & Rat-

cliff, 1992, 1995; McNamara, 1992, 1994). Our results, like those of McKoon and Ratcliff (1992), strongly suggest that they do not. The majority of our stimuli were not normative associates, and in the additional analyses, we demonstrated that the priming effects were not driven by the small number of items with high normative association values.

One might still argue, however, that although direct associations (as measured in a free association task) do not account for our results, mediated associations might. Researchers such as Balota and Lorch (1986) and McNamara and Altarriba (1988) have demonstrated priming effects in which *lion* primes *stripes*. They discussed these results in terms of spreading activation in an associative network, arguing that it is due to activation spreading from *lion* to *stripes* thorough the mediating node of *tiger*. McKoon and Ratcliff (1992) take a different approach, arguing that mediated priming is due to weak but direct co-occurrences within a moving window (as they operationalize it). Our discussion of event-based priming might at first glance appear to be an instance of mediated priming. For example, it could be suggested that *bakery* primes *bread* via the types of events that occur at bakeries. That is, the priming effects that we demonstrate could be viewed as mediated through a node that represents a specific event (a verb node, for example).

However, this is not the sort of mediation that we envision. We do not view our priming effects as being due to activation propagating along associative links in a localist lexical network. Nor is that mechanism required. BEAGLE, for example, correctly predicts the mediated priming effects of Balota and Lorch (1986) and McNamara and Altarriba (1988), as well as the long-lag mediated priming found by McNamara (1992).

The BEAGLE priming effects clearly cannot be due to activation spreading via a linked mediating node in the model, given that it contains no such connections. In the model, *lion* and *stripes* would become directly proximal to one another from occurrence in similar contexts, but they do not require a *tiger* node to mediate the spread of activation. Nor must these effects be due to direct co-occurrence – two words may become proximal in semantic space from occurrence in similar contexts, but do not necessarily need to directly co-occur.

Our idea is that the relations on which we have focused become instantiated directly in memory due to people's experience with the relevant types of events. *Bread* becomes related to *bakery* because bread is made there, and people go to bakeries to buy their bread. Hence when *bakery* is read, it activates *bread* due to this relational knowledge. This account differs crucially from the claim that there is a node that represents a specific event, with activation passing from *bakery* to that node, and then on to *bread*. In summary, the "mediated" aspect, if you will, of the priming effects demonstrated here occurs during learning.

7.2.3. Models and measures of semantic relatedness

Throughout this work we have discussed a variety of approaches used in determining semantic relatedness. Although there are similarities among them, there are also clear and important differences. Association norms, for example, are behavior, while co-occurrence is a fact about the statistical structure of the world, and hence something that people *learn* rather than something they do. Co-occurrence models are essentially statistical tabulations of that information in the world, which is the information from which human behavior is learned. Crucially (as examples from cognitive psychology have shown) the structure of behavior is not the same thing as the structure of the data that that leads one to produce that behavior. It would therefore be misguided to argue which measure is correct – we believe instead that the most informative approach is to determine what association is in the world, and from that produce a model that learns from these data and produces behavior as humans do.

Furthermore, there is more than one dimension to semantic relatedness. In recent work, Maki and Buchanan (2008) explored global similarities between measures of lexical semantic similarity. They compiled word–word similarities for concrete nouns and action verbs from featural representations (McRae et al., 2005; Vinson & Vigliocco, 2008), co-occurrence models (LSA, the Topic model, and BEAGLE), association norms (Nelson, McEvoy, & Schreiber, 1998), and word similarity computed from WordNet (Maki, McKinley, & Thompson, 2004). Using multidimensional scaling, hierarchical clustering, and factor analysis, they demonstrated that these measures of semantic similarity were separable, encoding somewhat different types of semantic similarity. Based on these results, Maki and Buchanan argue that co-occurrence models encode a distinct type of semantic similarity.

7.2.4. Association and higher-order relationships

Although we have argued against simple spreading activation accounts, our stimulus pairs are indeed associated if *association* is taken in its general sense of referring to the relationship between things and words that are contiguous in the world or in language. It would be quite odd indeed to argue that concepts such as *sale* and *shopper*, or *classroom* and *student* are not associated in this general sense. As discussed above, these patterns of co-occurrence do arise in language use, arguably because they reflect the concerns of language users. Interestingly, Prior and Bentin (2008) find that incidental associations are learned between word pairs that are presented and perceived as noun arguments in coherent sentences, but not when the sentences that contain them are semantically anomalous. The importance of this finding is that such associations are largely formed when words are encountered in meaningful contexts, where the two words form essential parts of an integrated scene or comprehensible event.

Learning – including the learning of event knowledge – involves the experience of repeated spatial and temporal contiguities of this sort. Importantly however, it also leads to emergent abstractions across these multiple experiences, so that the resultant knowledge reflects relationships among these emergent representations rather than undifferentiated lexical relationships. As Traxler et al. (2000) showed, these relationships reflect higher-order

information about event roles. We have presented similar findings in other work. In priming studies we found that verbs prime typical patients of the event the verb denotes (Ferretti et al., 2001), and that nouns prime verbs denoting events in which they are typical patients (McRae et al., 2005). In reading time studies, however, we have shown that these relationships are mediated by context, as would be expected if the relationship encoded more than simple co-occurrence relations. In single sentences, comprehenders show processing difficulty when an otherwise good patient for a verb (e.g. *saved money*) is made less likely because the combination of verb and agent earlier in the sentence (e.g. *the lifeguard saved*) evokes an event that decreases the likelihood of this patient (Bicknell, Elman, Hare, McRae, & Kutas, 2008). In addition, these sentences in turn are integrated into prior context if any is available, and a larger context can override the expectancies that the agent–verb combination might elicit in isolation (Race, Klein, Hare, & Tanenhaus, 2008). Thus context at various levels influences comprehenders' interpretation of the event being described and the roles that event participants play.

But although the larger pattern of results is inconsistent with simple normative association, this does not mean that knowledge of common events could not be represented in an associative network. Such networks are demonstrably capable of accounting for word–word priming effects of the type found here, and would offer a straightforward account of our finding that the number of intervening neighbors influences the degree of facilitation between word pairs. These networks implement one well-established account of meaningful relations between concepts (Anderson, 1983; Collins & Loftus, 1975), and indeed, the Rumelhart et al. (1986) model of a *room* schema, which captures crucial aspects of our account, is implemented as a spreading activation network. The important point here is that any spreading activation network that could account for the existing priming, sentence comprehension, and discourse results would need to include specific types of relations rather than simply undifferentiated association.

The NIN measure might be thought of as a proxy for the structure that a process model of semantic priming would operate upon. If a process such as spreading activation is viewed as energy diffusion through semantic space (because words in BEAGLE are not "connected" in the traditional semantic network sense), and words that exist in space on the way to the target from the prime may absorb a portion of this energy, then the number of intervening neighbors is a good approximation of the behavior this process would produce. If more words exist on the path of activation from *king* to *crown* than on the path from *crown* to *king*, facilitation is more likely to be seen in the latter case than the former. In addition, the density of the neighbors is an important factor: *n* intervening words close to the prime may be more likely to absorb energy than an equal number further away. A formal model of this process is beyond the scope of the current article, but the NIN measure can serve as a proxy for the behavior that might be produced by an attractor network in this space. Such networks have proven effective

at accounting for temporal dynamics of priming (Cree, McRae, & McNorgan, 1999; Mirman & Magnuson, 2008), but have not yet been adapted to a space such as that produced by BEAGLE.

8. Conclusion

Event-based knowledge and relations are playing an increasingly important role in a number of areas. They have been shown to influence sentence comprehension (Altmann & Kamide, 1999; Bicknell et al., 2008; Filik, 2008; Race et al., 2008; Trueswell, Tanenhaus, & Garnsey, 1994) and lexical ambiguity resolution (Vu et al., 2003). Event-based thematic relations have also been shown to influence performance on categorization and similarity rating tasks (Jones & Love, 2007; Lin & Murphy, 2001; Markman & Stilwell, 2001; Wisniewski & Bassok, 1999). Event knowledge, in the form of functional information regarding how people use objects, is implicated as a key component of understanding category-specific semantic deficits (Simmons & Barsalou, 2003; Warrington & Shallice, 1984). The present experiments contribute to our understanding of the specific types of relations that are encoded in semantic memory, are activated when words are read or heard, and influence language comprehension in general.

Acknowledgements

Appendix A. Event–people prime–target pairs for Experiment 1a

olympics	athlete
appointment	doctor
trial	judge
lecture	professor
baptism	priest
accident	policeman
war	soldier
wedding	bride
robbery	burglar
recess	kids
cruise	captain
reunion	friends
ballgame	players
anniversary	husband
birth	baby
divorce	wife
election	candidate
sale	shopper

Appendix B. Event-thing prime–target pairs for Experiment 1b

funeral	coffin
accident	cars
blizzard	snow
war	guns
parade	floats
regatta	boat
contest	prizes
eclipse	moon
graduation	diploma
sale	clothes
breakfast	eggs
election	ballot
banquet	food
baptism	water
dinner	plate
interrogation	table
recess	ball
marathon	shoes
picnic	blanket
olympics	medal
trip	luggage
surgery	scalpel
joust	sword
barbecue	hamburger
concert	speaker
robbery	money

Appendix C. Location–people and location–animal prime–target pairs for Experiment 2a

stable	horse
sandbox	children
hospital	doctor
river	fish
classroom	student
church	priest
palace	king
court	judge
office	secretary
airport	pilot
igloo	eskimo
laboratory	scientist
barn	cow
casino	gambler
rink	skater
theatre	actor
tavern	bartender
woods	deer
bakery	cashier
desert	camel
spa	masseuse
gym	athlete
beach	lifeguard
mall	shopper

Appendix D. Location–thing prime–target pairs for Experiment 2b

garage	car
pool	water
desert	sand
neighborhood	houses
meadow	grass
greenhouse	plants
laundromat	washer
airport	planes
casino	slots
bathroom	toilet
jungle	trees
cemetery	gravestones
closet	clothes
office	desk
barn	hay
factory	machines
warehouse	boxes
gym	weights
restaurant	tables
hotel	beds
sandbox	shovel
zoo	cages
bakery	bread
spa	jacuzzi
classroom	blackboard
salon	scissors
market	food
nightclub	liquor
alley	garbage
farm	tractor

Appendix E. Instrument–people prime–target pairs for Experiment 3a. The same pairs, in reversed order were used in Experiment 4

backpack	student
crayon	child
briefcase	businessman
balloon	clown
hose	gardener
saucepan	cook
detergent	mother
microscope	biologist
pen	teacher
wrench	plumber
saddle	rider
broom	janitor
knife	chef
rod	fisherman
compass	hiker
lantern	camper
razor	barber
needle	seamstress
rag	maid
shampoo	hairdresser
stopwatch	coach
ladder	painter
helmet	motorcyclist
cards	gambler

Appendix F. Instrument–thing prime–target pairs for Experiment 3b

battery	car
beaker	liquids
bowl	cereal
briefcase	papers
broom	floor
crayon	pictures
detergent	clothes
fridge	food
furnace	house
hoe	garden
key	door
ladle	soup
mug	coffee
net	fish
oven	cookies
pole	flag
pump	tires
razor	face
scissors	hair
screen	bugs
spatula	pancakes
strainer	water
tweezers	eyebrow
wrench	bolt

References

Altmann, G. T. M. (1999). Thematic role assignment in context. *Journal of Memory and Language, 41*, 124–145.

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition, 73*, 247–264.

Anaki, D., & Henik, A. (2003). Is there a "strength effect" in automatic semantic priming? *Memory and Cognition, 31*, 262–272.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 336–345.

Bicknell, K., Elman, J. L., Hare, M., McRae, K., & Kutas, M. (2008). Online expectations for verbal arguments conditional on event knowledge. In: Love, B. C., McRae, K., & Sloutsky, V. M. (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 2220–2225). Austin, TX: Cognitive Science Society.

Boland, J., Tanenhaus, M. K., Garnsey, S., & Carlson, G. (1995). Verb argument structure in parsing and interpretation: Evidence from wh-questions. *Journal of Memory and Language, 34*, 774–806.

Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory. In Dietrich, E., & Markman, A. B., (Eds.), *Cognitive dynamics: Conceptual change in humans and machines*.

Camblin, C., Gordon, P., & Swaab, T. (2007). The interplay of discourse congruence and lexical association during sentence processing: Evidence from ERPs and eye tracking. *Journal of Memory and Language, 56*, 103–128.

Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments and Computers, 25*, 257–271.

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

Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science, 23*, 371–414.

de Groot, A. M. B. (1984). Primed lexical decision: Combined effects of the proportion of related prime–target pairs and the stimulus-onset asynchrony of prime and target. *Quarterly Journal of Experimental Psychology, 36A*, 253–280.

den Heyer, K., Briand, K., & Dannenbring, G. (1983). Strategic factors in a lexical decision task: Evidence for automatic and attention-driven processes. *Memory & Cognition, 11*, 374–381.

Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language, 44*, 516–547.

Filik, R. (2008). Contextual override of pragmatic anomalies: Evidence from eye movements. *Cognition, 106*, 1038–1046.

Fischler, I. (1977). Semantic facilitation without association in a lexical decision task. *Memory & Cognition, 5*, 335–339.

Garrod, S., & Terras, M. (2000). The contribution of lexical and situational knowledge to resolving discourse roles: Bonding and resolution. *Journal of Memory and Language, 42*, 526–544.

Hare, M., Elman, J. L., Tabaczynski, T., & McRae, K. (in press). The wind chilled the spectators, but the wine just chilled. *Cognitive Science*.

Hess, D., Foss, D., & Carroll, P. (1995). Effects of global and local context on lexical processing during language comprehension. *Journal of Experimental Psychology: General, 124*, 62–82.

Hutchinson, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review, 10*(4), 785–813.

Jones, M. N., & Kintsch, W. (2006). Asymmetric similarity in a self-organizing lexicon. In *Poster presented at the 47th meeting of the psychonomic society*. Houston, Texas.

Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language, 55*, 534–552.

Jones, M., & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology, 55*, 196–231.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review, 114*, 1–37.

Kako, E. (2006). Thematic role properties of subjects and objects. *Cognition, 101*, 1–42.

Lancaster, J. S., & Barsalou, L. W. (1997). Multiple organizations of events in memory. *Memory, 5*, 569–599.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: the Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211–240.

Lin, E., & Murphy, G. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General, 130*, 3–28.

Lupker, S. J. (1984). Semantic priming without association: A second look. *Journal of Verbal Learning and Verbal Behavior, 23*, 709–733.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes, 9*, 157–201.

Maki, W. S., & Buchanan, E. (2008). Latent structure in measures of associative, semantic, and thematic knowledge. *Psychonomic Bulletin & Review, 15*, 598–603.

Maki, W. S., McKinley, L. N., & Thompson, A. G. (2004). Semantic distance norms compued from and electronic dictionary (Word-Net). *Behavior Research Methods, Instruments, & Computers, 36*, 421–431.

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence, 13*, 329–358.

McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 1155–1172.

McKoon, G., & Ratcliff, R. (1995). Sequential effects in lexical decision: Tests of compound-cue retrieval theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1380–1388.

McKoon, G., & Ratcliff, R. (2005). Interactions of meaning and syntax: Implications for models of sentence comprehension. *Journal of Memory and Language, 56*, 270–290.

McNamara, T. P. (1992). Priming and theories of memory: A reply to Ratcliff and McKoon. *Psychological Review, 101*, 185–187.

McNamara, T. P., & Altarriba, J. (1988). Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language, 27*, 545–559.

McNamara, T. P. (2005). *Semantic priming: Perspectives from memory and word recognition*. New York, NY: Psychology Press.

McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 558–572.

McRae, K., Hare, M., Elman, J. L., & Ferretti, T. R. (2005). A basis for generating expectancies for verbs from nouns. *Memory and Cognition, 33*, 1174–1184.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language, 38*, 283–312.

Mirman, D., & Magnuson, J. S. (2008). Attractor dynamics and semantic neighborhood density: Processing is slowed by near neighbors and speeded by distant neighbors. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 65–79.

Morris, R. K. (1994). Lexical and message-level sentence context effects on fixation times in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 92–103.

Morrow, D. G., Greenspan, S. L., & Bower, G. H. (1987). Accessibility and situation models in narrative comprehension. *Journal of Memory and Language, 26*, 165–187.

Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 863–883.

Myers, J. L. (1979). *Fundamentals of experimental design*. Boston, MA: Allyn & Bacon.

Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*, 226–254.

Neely, J. H., Keefe, D. E., & Ross, K. (1989). Semantic priming in the lexical decision task: Roles of prospective prime-generate expectancies and retrospective semantic matching. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 6*, 1003–1019.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. <http://www.usf.edu/FreeAssociation/>.

Pollatsek, A., & Well, A. D. (1995). On the use of counterbalanced designs in cognitive research: A suggestion for a better and more powerful analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 785–794.

Prior, A., & Bentin, S. (2008). Word associations are formed incidentally during sentential semantic integration. *Acta Psychologica, 127*, 57–71.

Race, D., Klein, N., Hare, M., & Tanenhaus, M. K. (2008). What do shoppers expect to save? Agents influence patients via events, not associative priming. In *Poster presented at the 18th annual conference on human sentence processing*. NC: Chapel Hill.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In D. E. Rumelhart & J. L. McClelland (Eds.). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 7–57). Cambridge, MA: MIT Press.

Sahlgren, M. (2005). An introduction to random indexing. In *Proceedings of the methods and applications of semantic indexing workshop at the 7th international conference on terminology and knowledge engineering*. TKE 2005.

Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory and Cognition, 18*, 1191–1210.

Simmons, W. K., & Barsalou, L. W. (2003). The similarity-in-typography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology, 20*, 451–486.

Steyvers, M. (2000). Modeling semantic and orthographic similarity effects on memory for individual words. Indiana University, unpublished doctoral dissertation.

Stowe, L. (1989). Thematic structures and sentence comprehension. In G. N. Carlson & M. K. Tanenhaus (Eds.), *Linguistic structure in language processing*. Dordrecht: Kluwer.

Tanenhaus, M. K., & Carlson, G. N. (1989). Lexical structure and language comprehension. In W. Marslen-Wilson (Ed.), *Lexical representation and process* (pp. 529–561). Cambridge: The MIT Press.

Tanenhaus, M. K., Carlson, G. N., & Trueswell, J. (1989). The role of thematic structures in interpretation and parsing. *Language and Cognitive Processes, 4*, SI 211–SI 234.

Traxler, M. J., Foss, D. J., Seely, R. E., Kaup, B., & Morris, R. K. (2000). Priming in sentence processing: Intralexical spreading activation, schemas, and situation models. *Journal of Psycholinguistic Research, 29*, 581–595.

Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language, 33*, 285–318.

Van Berkum, J. J. A., Brown, C. M., Zwisterlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating upcoming words in discourse: Evidence from ERPs and reading times. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 443–467.

Van Berkum, J. J. A., Hagoort, P., & Brown, C. M. (1999). Semantic integration in sentences and discourse: Evidence from the N400. *Journal of Cognitive Neuroscience, 11*, 657–671.

Vinson, D., & Vigliocco, G. (2008). Semantic feature norms for a large set of objects and events. *Behavior Research Methods, Instruments, & Computers, 40*, 183–190.

Vu, H., Kellas, G., Petersen, E., & Metcalf, K. (2003). Situation-evoking stimuli, domain of reference, and the incremental interpretation of lexical ambiguity. *Memory & Cognition, 31*, 1302–1315.

Warrington, E., & Shallice, T. (1984). Category specific semantic impairments. *Brain, 107*, 829–854.

Wisniewski, E. J., & Bassok, M. (1999). Stimulus compatibility with comparison and integration. *Cognitive Psychology, 39*, 208–238.

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