Cognition 111 (2009) 151–167

Contents lists available at ScienceDirect

Cognition

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

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

Article history:
Received 15 January 2007
Revised 12 January 2009
Accepted 30 January 2009

Keywords:
Semantic priming
Association
Sentence processing
Event representation
Language comprehension
Semantic memory

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 relations, 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).

0010-0277/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2009.01.009

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