

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Lexical Ambiguity Resolution in a Constraint Satisfaction Network

Permalink

<https://escholarship.org/uc/item/9277h0zd>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 11(0)

Author

Masson, Michael E. J.

Publication Date

1989

Peer reviewed

Lexical Ambiguity Resolution in a Constraint Satisfaction Network

Michael E. J. Masson

University of Victoria

ABSTRACT

Behavioral evidence supports the claim that in the absence of a strongly biasing context multiple meanings of an ambiguous word are activated, particularly when the two meanings occur with equal frequency. A simple constraint satisfaction system, based on a Hopfield network and incorporating a distributed memory scheme, is shown to account for results from a cross modal priming paradigm typically interpreted as evidence for multiple access. The model demonstrates that the power of an ambiguous word to facilitate identification of targets related to either of its two meanings may be produced by *selective* activation of just one meaning. Selective activation is driven by simultaneous processing of the ambiguous prime and the associated target word, with the unambiguous target determining the appropriate interpretation of the prime. The model also provides the basis for a reinterpretation of a number of other empirical results concerning lexical ambiguity resolution.

INTRODUCTION

The identification of a word that has at least two clearly distinct interpretations presents a serious challenge to models of word identification. The critical issue is whether all or only one of an ambiguous word's meanings are activated upon presentation of the word. Although certain kinds of contextual information apparently produce a form of selective access of a homograph's lexical entries (Duffy, Morris, & Rayner, 1988; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Tabossi, 1988; Van Petten & Kutas, 1987), there is evidence to suggest that multiple access occurs in the absence of biasing context. Using a cross modal priming paradigm and balanced homographs (items with two equally frequent meanings), Seidenberg et al. (1982) found that hearing a homograph at the end of a nonbiasing sentence context produced equal priming of visual targets related to either of the homograph's two interpretations. Similarly, Duffy et al. (1988) found longer gaze durations, relative to control words, on balanced homographs that appeared following a nonbiasing phrase. The additional processing time was attributed to required selection of one of the two activated entries for integration with the context.

Simultaneous activation of multiple lexical entries is fully compatible with the classic semantic network view of the lexicon (e.g., Collins & Loftus, 1975; Collins & Quillian, 1972). This view holds that each lexical entry is represented as a unique node in the network, and it is quite possible for multiple nodes to be highly active. In fact, priming effects are assumed to be the result of activation spread across links between nodes. But strong, simultaneous activation of concepts, particularly unrelated concepts as in the case of distinct meanings of an ambiguous word, is problematic for models based on a distributed memory scheme (e.g., Eich, 1985; McClelland & Rumelhart, 1985). In a distributed memory system a concept is expressed as a pattern of activation across an entire network of nodes. Only one pattern of activation can be instantiated at any moment, therefore it would be impossible for two unrelated concepts (representing the two different meanings of an ambiguous word) to simultaneously dominate the network's pattern of activation. For this reason, demonstrations of multiple access constitute a serious challenge to this class of models.

Kawamoto and Anderson (1984) developed a neural network model incorporating a distributed memory representation that produced simultaneous activation of multiple meanings of an ambiguous word. In keeping with the limitations on multiple activation imposed by a distributed

MASSON

memory scheme, however, their reported simulations included no instance of two concepts simultaneously achieving activation levels greater than .5, where asymptotic activation was 1.0. Although this could be considered a case of multiple activation, it is not clear from their simulations whether this rather low level of activation would be adequate to produce priming effects equal in magnitude to those obtained with an unambiguous word (Seidenberg et al., 1982).

A CONSTRAINT SATISFACTION NETWORK MODEL

I have attempted to apply a simple constraint satisfaction model, based on a Hopfield network (Hopfield, 1982; Hopfield & Tank, 1986), to the problem of simulating cross modal priming results that have been taken as evidence for multiple access in the lexicon (Seidenberg et al., 1982). The network consists of a group of binary valued nodes, and concepts are represented as unique patterns of on/off states across these nodes. Only one pattern of activation can be represented in the system at a given instant, implying that the system will not allow multiple concepts to be fully activated simultaneously. But this restriction does not mean that multiple concepts cannot be partially activated at the same time. Degree of activation of a concept in this system is conceived as the proportion of nodes in the currently instantiated pattern that match the pattern defining the concept. When a pattern of activation is instantiated in the network it may partially match a number of different concepts. The maximum amount of simultaneous activation of a set of concepts, however, is limited by the similarity of their defining patterns of activation.

The system's knowledge about and potential for instantiating different concepts lies in the strengths of connections between nodes. New concepts are acquired by adjusting the connection weights according to a prescription of the general type proposed by Hebb (1949). The weight of the link connecting any two nodes is altered as a function of the on/off states adopted by the two nodes when a new concept is instantiated. Specifically, if both nodes are in the same state (both on or both off) then the link between them is increased in strength, but if the two nodes are in different states (one on and one off) their link is reduced:

$$\Delta w_{ij} = n_i n_j,$$

where w_{ij} represents the weight of the link between nodes i and j , and n_i and n_j represent the states taken by those nodes when the new concept is activated. A node in the on state takes the value 1 and a node in the off state is assigned the value -1.

Pattern Recognition

Simulation of word identification is treated as a pattern completion problem in the model. Once the system has encoded a set of concepts, it can be provided a pattern of activation that partially matches a target concept. The system can then gradually change the pattern of activation until it reconstructs the entire pattern of the target concept. This is accomplished through a process of asynchronous updating of each node in the network. If the net activation sent to a node from all other nodes in the network exceeds some threshold, the node is set to the *on* state, otherwise it is set to the *off* state. The activation received by a node is a function of the on/off states of the other nodes in the system and the weight of the links between those nodes and the node selected for updating:

$$a_i = \sum_{j \neq i} w_{ij} n_j,$$

where a_i represents the amount of activation directed to the selected node. In the simulations reported here the threshold for setting a node to the *on* state was zero. This updating scheme acts

MASSON

as a constraint satisfaction network, driving the network to a pattern of activation that represents a local maximum goodness of fit, and is stable in the sense that any further updates will not produce adjustments in the pattern. Under appropriate circumstances (e.g., the concepts learned by the network are few in number relative to the number of nodes in the network, and are not highly similar to one another), the stable pattern that is achieved will correspond to one of the concepts learned by the system (Hopfield, 1982).

To simulate word identification, a network consisting of two sets of nodes was constructed. One set of nodes represents the perceptual (visual and auditory) characteristics of a word, and another set represents its conceptual features. Each perceptual node is linked to each conceptual node, and the conceptual nodes are fully interconnected as well. Upon presentation of a word, perceptual input is immediately read into the appropriate subset of perceptual nodes (e.g., visual) and the other subset of perceptual nodes (auditory) are set to zero, indicating no relevant input is present in the other modality. No attempt is made in the current version of the model to simulate interactive effects in perception. Asynchronous updating is applied to the conceptual nodes, which begin in either a random state or a pattern determined by a previously presented stimulus. The perceptual nodes continue to hold the pattern of activation dictated by the sensory input. The updating scheme drives the pattern of activation among the conceptual nodes from the initial state to a stable pattern representing the concept designated by the perceptual input. Reaching a learned stable state constitutes full activation and identification of the target concept, and the number of updating cycles required to move the network to the stable state is taken as a measure of word identification time.

SIMULATION OF LEXICAL AMBIGUITY RESOLUTION

The critical issue was whether a system that inherently prohibits strong, simultaneous activation of unrelated concepts could simulate cross modal priming results that have been taken as evidence for multiple access of an ambiguous word's meanings. An attempt was made to simulate the results of Experiment 1 reported by Seidenberg et al. (1982). In this study subjects heard neutral phrases ending with an ambiguous or unambiguous prime word, then named aloud a visually presented target (e.g., *Joe buys the straw-HAY*). Presentation of the target occurred either immediately after the prime was pronounced or 200 msec later. In the case of an ambiguous prime the target was related to one of its meanings and when an unambiguous prime was used the target was either related or unrelated to it. When the target appeared immediately after the prime, homograph primes produced facilitation in naming equal to that produced by unambiguous primes. After a 200-msec delay, however, greater facilitation was obtained with unambiguous primes. Seidenberg et al. concluded that immediately after hearing an ambiguous prime both meanings were activated, thereby producing facilitation to either target. But after a delay one interpretation was arbitrarily selected so only one of the targets would be facilitated and the other would not. The average facilitation, then, would be smaller than that obtained with the unambiguous prime.

Simulation of these results involved a version of the network described earlier, consisting of 45 auditory, 40 visual, and 40 conceptual nodes. An ambiguous word was simulated by constructing two items with identical patterns of activation among perceptual nodes, and unrelated patterns in the conceptual nodes. These two items were encoded using the Hebbian learning rule. In addition, two unambiguous words related to each meaning of the homograph were also encoded. The patterns of activation in the conceptual nodes of related words overlapped in 29 of the 40 nodes. These items were used to simulate the three priming conditions tested by Seidenberg et al. (1982). The homograph and one of the unambiguous words related to each of its meanings served as primes and the other unambiguous words were targets. It was assumed that information in the neutral phrase would leave the conceptual nodes in a random pattern of activation because none of

MASSON

TABLE 1

MEAN CYCLES REQUIRED TO IDENTIFY A VISUAL TARGET IN THE
SIMULATION OF THE CROSS MODAL PRIMING PARADIGM

Target delay	Prime		
	Related ambiguous	Related unambiguous	Unrelated unambiguous
0 msec	174.8	175.2	208.4
200 msec	140.7	128.4	167.5

its constituents were strongly related to the primes or targets. Presentation of the prime was simulated by loading its auditory pattern into the network and onset of the visual target was simulated by loading the appropriate visual pattern into the network.

In the immediate condition it was assumed that some minimal processing of the prime would occur as it was being pronounced and that additional processing would continue during the early stages of target processing. Therefore the network was updated for 10 cycles using only the prime's auditory pattern, with the visual nodes set to zero. Then the visual pattern for the target was loaded and the network was updated under the influence of both perceptual patterns for a further 65 cycles. Finally the auditory pattern was set to zero to reflect selective processing of the visual target and updates continued until the system reached a stable state or 400 updates had occurred since target onset. In the 200-msec delay condition, the prime's auditory pattern was allowed to influence the system for 75 cycles at which point the target's visual pattern was loaded and the auditory pattern was set to zero. The system then continued to update until the target was identified or the system ran for 400 update cycles since target onset. This sequence was formulated on the assumption that in the delay condition processing of the auditory prime was terminated by onset of the visual target.

The simulation was run 100 times for each of the two delay conditions, with each run based on a set of words randomly selected within the constraints described earlier. Each run involved 10 trials in each priming condition. Despite the high degree of similarity among the concepts encoded by the system, only 13% of the trials failed to move into the stable state that defined the target word. The mean number of cycles required to identify the target as a function target delay and type of prime is shown in Table 1. These results very closely approximate the data reported by Seidenberg et al. (1982), with the exception that the simulation produced shorter response times in the delay condition. Seidenberg et al. found longer response times after a 200-msec delay in their first experiment, but the delay variable was manipulated between subjects and a within-subject manipulation would be needed to obtain a more reliable assessment of the overall effect of delay on response time. Moreover, other experiments they reported are consistent with the simulation, as are independent replications of their experiment (e.g., Van Petten & Kutas, 1987). The model produces faster identification in the delay condition because the target's visual pattern unilaterally influences activation in the conceptual nodes as soon as it appears. There is no competition from the auditory prime trying to move the system to a slightly different activation pattern in the conceptual nodes.

The model's success in producing equal facilitation with ambiguous and unambiguous primes in

MASSON

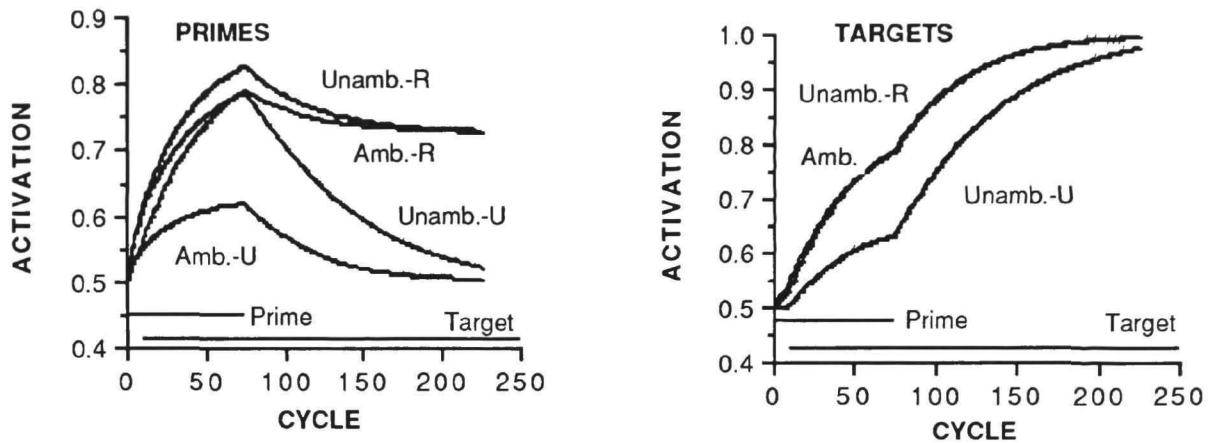


FIGURE 1. ACTIVATION VALUES FOR PRIMES (LEFT PANEL) AND TARGETS (RIGHT PANEL) IN THE IMMEDIATE CONDITION OF THE CROSS MODAL PRIMING PARADIGM. AMBIGUOUS AND UNAMBIGUOUS PRIMES WERE EITHER RELATED (R) OR UNRELATED (U) TO THE TARGETS. HORIZONTAL LINES INDICATE TIME COURSE OF SENSORY INPUT FROM PRIME AND TARGET.

the immediate condition was not due to multiple activation of both meanings of an ambiguous prime. Rather, simultaneous processing of the auditory prime and visual target in the immediate condition allowed the target to drive the system toward the relevant interpretation of the ambiguous prime. This process can be seen in Figure 1, which represents the activation values for prime (left panel) and target (right panel) concepts in the immediate condition. The activation value for a concept at a given instant is the proportion of the network's conceptual nodes that currently match the pattern of on/off states corresponding to that concept. A value of .5 indicates that a concept is at resting level in the sense that a concept would be expected to have half its nodes in common with a randomly chosen pattern of activation. The behavior shown in Figure 1 is not an example of *backward* priming, where the target activates a related meaning of the ambiguous prime, then the selected meaning sends activation back to the target (Glucksberg, Kreuz, & Rho, 1986). The mutually supportive activation of the conceptual nodes produced by simultaneous processing of the ambiguous prime and the target is more closely related to what Van Petten and Kutas (1987) referred to as *mutual* priming. In a distributed memory system this kind of interaction between closely coupled inputs involves simultaneous influence over the entire network of nodes. The influence is symbiotic when the two inputs represent conceptually related items. Equal priming by ambiguous and unambiguous primes occurred despite slightly lower activation of the relevant interpretation of the ambiguous prime. This is because patterns of activation representing homographs generally have higher goodness of fit values than unambiguous items, giving them greater power to influence the direction of change in the network's pattern of activation. The goodness of fit advantage for homographs is a product of the two encodings of their perceptual patterns (once for each meaning of a homograph) compared to only one encoding of the perceptual pattern for an unambiguous word.

When the target was delayed the system had no basis for selecting the appropriate interpretation and was therefore incorrect in its selection on about half of the trials. In those instances the auditory prime moved the conceptual nodes toward a pattern of activation irrelevant to the target. Consequently, the average level of activation of relevant and irrelevant interpretations of the ambiguous primes were similar and lower than activation of the unambiguous primes, as seen in the left panel of Figure 2. Moreover, the ambiguous prime was effective in facilitating target identification only half the time, that is, on those occasions when the relevant interpretation

MASSON

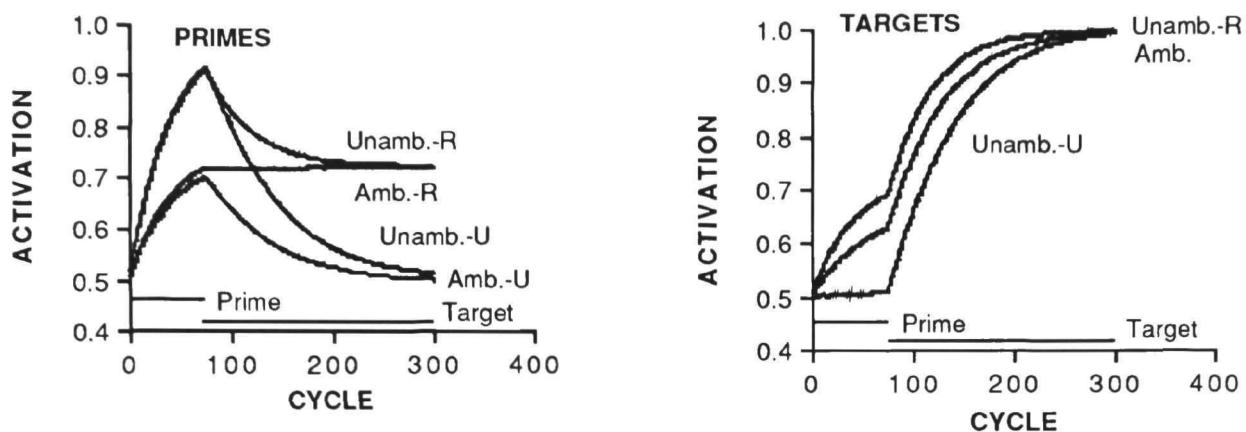


FIGURE 2. ACTIVATION VALUES FOR PRIMES (LEFT PANEL) AND TARGETS (RIGHT PANEL) IN THE DELAYED CONDITION OF THE CROSS MODAL PRIMING PARADIGM. AMBIGUOUS AND UNAMBIGUOUS PRIMES WERE EITHER RELATED (R) OR UNRELATED (U) TO THE TARGETS. HORIZONTAL LINES INDICATE TIME COURSE OF SENSORY INPUT FROM PRIME AND TARGET.

happened to be selected. The average growth of activation values of target items is shown in the right panel of Figure 2 and reflects the differences in observed facilitation effects.

Other Empirical Results

The model could be extended to account for another effect reported by Seidenberg et al. (1982). Using a biased context, they found evidence for selective activation of one interpretation of an ambiguous prime even when the target followed immediately after the prime. The context phrase contained a word strongly related to one interpretation of the ambiguous prime (e.g., *Although the farmer bought the straw-HAY*). Context effects of this form could be simulated by assuming that context words are identified without driving the conceptual nodes completely into the appropriate pattern of activation. Rather, an approximation to the known pattern would be sufficient to identify the word. This would allow parts of the pattern established by earlier context words to survive in the network until the ambiguous prime was presented. At that point the system would be in a state that slightly favored the relevant interpretation and that is the meaning that likely would be selected. The model could similarly account for results such as those obtained by Duffy et al. (1988) and Tabossi (1988) in which contexts not containing words directly associated with an ambiguous word nevertheless influenced which interpretation was selected. For example, context could activate elements of the conceptual nodes relevant to a salient feature of one meaning of the homograph causing that meaning to be selected when the homograph was presented (Tabossi, 1988).

In their study of event related potentials, Van Petten and Kutas (1987) replicated the Seidenberg et al. (1982) study except that contexts and targets both were visually presented and the contexts contained information to bias one interpretation of the homograph prime. The interesting result was that when a target related to the unbiased interpretation was presented immediately after termination of the homograph prime, event related potentials indicated that the target was initially responded to as though it were a contextually unrelated word. But 500 msec after target onset cortical activity suggested that the alternative interpretation of the homograph had been instantiated. This is the result that would be expected of the constraint satisfaction network model proposed here. Presentation of the target would eventually change the interpretation applied to the homograph prime, overriding the early influence of the preceding context.

MASSON

There is a second kind of evidence that supports the multiple access view of lexical ambiguity resolution. In their study of eye movements during reading of ambiguous words, Duffy et al. (1988) found that balanced homographs were viewed an average of 18 msec longer than control words, provided that the preceding context was neutral with respect to the homograph's interpretation. Their explanation was that both meanings had been activated and a time consuming selection process was required in order for the word to be integrated with ongoing comprehension processes. But the viewing time effect could have resulted from an important difference between the homograph and control words. Duffy et al. equated these items for frequency, but they did so on the basis of occurrences of tokens rather than types. Thus for a balanced homograph, each of its meanings would be experienced in a relevant context by an average reader only half as many times as the control word. The longer viewing times for homographs may actually reflect the lower frequency of the arbitrarily selected interpretation of the homograph, rather than activation of both interpretations.

CONCLUSION

Results from the cross modal priming paradigm have significantly influenced theories of lexical ambiguity resolution and, more generally, theories of lexical access and language comprehension. These results have encouraged the assumption that multiple meanings of ambiguous words are initially activated then a context appropriate sense is selected (e.g., Kintsch, 1988; Seidenberg, 1985). The constraint satisfaction network model described here, however, was able to reproduce the results on which this assumption was founded even though it clearly involved selective access of lexical ambiguity. Although this demonstration does not prove that lexical ambiguities are processed by selective activation, it does show that a multiple access interpretation is not dictated by the available empirical evidence, and that a selective access account of the data is plausible. In addition, the model represents one example of how distributed memory systems can account for data that imply simultaneous instantiation of incompatible patterns of activation.

The present version of the model makes no claims about the key theoretical idea behind the multiple access proposal: modularity of the lexicon. A modular view of the lexicon assumes that lexical activation is independent of contextual information provided by sentence comprehension (Fodor, 1983; Seidenberg, 1985). Therefore, identification of an ambiguous word should involve multiple activation of its meanings unless *lexical* processes influence selection of one meaning (Seidenberg et al., 1982). In the simulations reported here, lexical activity was influenced only by sensory input directly impinging on the lexicon or by activity within the lexicon itself. No influence from contextual information based on sentence comprehension was incorporated and in this sense the model preserves the assumption of modularity in the lexicon. The fact that this system was able to simulate relevant empirical results while engaging in selective activation suggests an uncoupling of the issues of modularity and lexical ambiguity resolution.

REFERENCES

- Collins, A. M., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review*, *82*, 407-428.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*, 240-248.
- Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, *27*, 429-446.
- Eich, J. M. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review*, *92*, 1-38.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

MASSON

- Glucksberg, S., Kreuz, R. J., & Rho, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12**, 323-335.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science*, **79**, 2554-2558.
- Hopfield, J. J., & Tank, D. W. (1986). Computing with neural circuits: A model. *Science*, **233**, 625-633.
- Kawamoto, A. H., & Anderson, J. A. (1984). Lexical access using a neural network. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society* (pp. 204-213).
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, **95**, 163-182.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, **114**, 159-188.
- Seidenberg, M. S. (1985). The time course of information activation and utilization in visual word recognition. In D. Besner, T. G. Waller, & G. E. MacKinnon (Eds.), *Reading research: Advances in theory and practice* (Vol. 5, pp. 199-252). New York: Academic Press.
- Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge toward processing. *Cognitive Psychology*, **14**, 489-537.
- Tabossi, P. (1988). Accessing lexical ambiguity in different types of sentential contexts. *Journal of Memory and Language*, **27**, 324-340.
- Van Petten, C., & Kutas, M. (1987). Ambiguous words in context: An event-related potential analysis of the time course of meaning activation. *Journal of Memory and Language*, **26**, 188-208.

ACKNOWLEDGEMENT

This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada and by computing facilities provided by the University of Victoria.