

A MODEL OF LEXICAL ACCESS OF AMBIGUOUS WORDS

Garrison W. Cottrell
Dept. of Computer Science
University of Rochester
Rochester, N.Y. 14627

ABSTRACT

Recent psycholinguistic work in the study of lexical access has supported a modular view of the process. That is, lexical access proceeds independently of the sentential context. Herein we describe a connectionist model of the process which retains modularity, explains apparent anomalies in the results, and makes empirically verifiable predictions.

INTRODUCTION

Within the domain of Artificial Intelligence there is considerable interest in parallel architectures for both machines and computational models (cf. Lesser & Erman 1977, Fahlman, 1980; Hillis, 1981; Fahlman, Hinton & Sejnowski, 1983) in part because of their promise as avenues for solving the fundamental AI problem of *search*. One such computational paradigm which has met with considerable enthusiasm and skepticism is the connectionist approach developed by Feldman and Ballard (1982; Feldman 1982). The theory was developed to reflect the current understanding of the information processing capabilities of neurons, and consequently the type of processing it supports is of a spreading activation/mutual inhibition character. While it at an early stage, the paradigm has been successfully applied in models of visual recognition of noisy inputs (Sabbah, 1982), motor control (Addanki, 1983), limited inference in semantic networks (Shastri & Feldman, 1984) and word sense disambiguation (Cottrell & Small, 1983). We believe it can be a useful cognitive modelling tool as well. Our intention here is to demonstrate how a simple connectionist model of a low level process in sentence comprehension can be effective in explaining psychological results.

We have two goals in building a cognitive model: To explain the existing data and to make empirically verifiable predictions. On the first point, in order to explain the data, it must be possible to form a clear correspondence between the elements of the model and the elements of the world that the theory attempts to explain. While many would argue, and have (Feigenbaum & Feldman 1963), that the neuronal level is the wrong place to start on such an enterprise, we claim that in order to explain the wealth of psycholinguistic

data on low level language processing*, the correspondence must be at a level below the functional; that the *mechanisms* involved in carrying out these functions must be considered if we are ever to have real explanatory power. Functional level models are effective in demonstrating what functions must be carried out; mechanism level models are better at explaining data from processing tasks. Second, we share the goal of all cognitive modelers to make predictions. Without this, there is no way to falsify the claims of the model. The model we present of lexical access makes several predictions which may be falsified or substantiated by subsequent research.

Within the domain of sentence processing various levels of analysis have been identified and generally agreed upon (eg. phonological, lexical, syntactic, semantic and pragmatic). However, the question of the characteristics of the interaction between these levels has been the focus of much study and debate (Forster, 1979; Marslen-Wilson & Tyler, 1980). The question is whether these systems can be regarded as independent modules or whether processing at one level can influence processing at another. The process of lexical access (defined below) presents a unique opportunity for modeling. The lexical level of processing has been intensively studied in recent years by psycholinguists who have focused precisely on the modularity question. What has emerged is a fairly well understood set of results which appears to resolve the question in favor of the modular view. Apart from the obvious consequences of the modularity issue for their field, researchers in AI should also be interested in this because obtaining the correct meaning of a word from the lexicon represents a search problem, due to large number of meanings of the most frequently used words. We discuss this research and then present a simple model of the lexical access process which explains apparent anomalies in the psycholinguistic results and makes empirically verifiable predictions.

*By "low level" we do not mean to imply the functional level/mechanism level distinction. We simply mean early stages of processing, such as phonological and lexical.

LEXICAL ACCESS

The process of accessing all of the information about a word, phonological codes, orthographic codes, meaning and syntactic features is called *lexical access*. We will mainly be concerned here with the access of meaning and syntactic class, and will use the term "lexical access" to refer to this process. It is useful to distinguish three stages the processing of lexical items, of which access is the second: (1) decoding the input and matching it with a lexical item, (2) accessing the information about that item, and (3) integrating that information with the preceding context. These are termed prelexical, lexical and postlexical processing, respectively. An important research question is discovering whether, to what degree, and through what channels these levels interact. Does each level only receive the completed output of the previous level (the "modular" view), or can processing at one level affect processing at adjacent or even more distant levels (the "interactive" view), or is the answer somewhere between these extremes?

Recent studies in lexical access have borne directly on the question of whether preceding context only has influence at the integration (postlexical) level or whether it can affect the lexical access level. A common tool in this research is to use ambiguous words and to study the effects of context on the processing of such words. The empirical question is whether the context of a sentence constrains the search in the lexicon for the contextually appropriate meaning of a word or not. The interactive view holds that context affects the lexical access level, so that only a single meaning is accessed (the Prior Decision Hypothesis). The modular view holds that all meanings of the word are initially accessed, since the lexical access mechanism can't "know" what the context requires, and all meanings are then passed to the integration level, where context selects the proper one (the Post Decision Hypothesis). Early research produced mixed results, some studies supporting one hypothesis, some the other (Conrad, 1974; Foss and Jenkins, 1973; Holmes, 1977; Lackner and Garret, 1972; Swinney and Hakes, 1976). However, most of these studies only looked at one time point of the process, which as later results show, explains the discrepancy.

Recent work by Swinney (1979) and others (Tanenhaus, Leiman, and Seidenberg, 1979; Seidenberg, Tanenhaus, Leiman, and Bienkowski, 1982) using the semantic priming measure has shown that the *time course* of lexical access is important. (We will only discuss the latter experiments here, referred to hereafter as STL.B. This study is the most comprehensive to date.) They use the *semantic priming* effect to measure the activation of a meaning of a word. People are faster and more accurate on various word tasks if the word is preceded by an associatively or semantically related word (see Meyer & Schvaneveldt, 1971). Semantic priming has also been shown to work cross-modally, i.e., a spoken word can

prime a visually presented word (Swinney, *et. al.*, 1979). In the STL.B study, the subject listened to a sentence containing an ambiguous priming word while being required to say a word (the *target*) flashed on a screen. By using targets semantically related to one meaning of the priming word (with appropriate controls) they were able to test the activation of the different meanings of the ambiguous word at different time points in the sentence. When the target is immediately following the prime, STL.B found priming from both meanings, (with one exception, discussed below) but when the target is 200 milliseconds later, only priming from the appropriate meaning for the sentence is found. This is interpreted as evidence that people initially access all meanings of a word followed by rapid decision (for sentence final words).

In addition finding a narrow decision window, STL.B discussed two types of context which differ in their effects on lexical access. They contrasted *pragmatic* context, resulting from world knowledge with *semantic* (or priming) context, resulting from associative and semantic relationships between word meanings, as in the following sentences.

- (1) The man walked on the *deck*.
- (2) The man inspected the ship's *deck*.
- (3) The man walked on the ship's *deck*.

The first sentence contains a pragmatic bias towards the "ship" related meaning of deck; one is more likely to walk on that kind. The second sentence contains a word highly semantically related to one meaning (ship -> deck). The third contains both types of information. They did experiments which contained a completely neutral context, a pragmatic context, or a semantic context. The results were that multiple access was obtained for neutral and pragmatic context, but selective access (only one reading active at the end of the word) for the semantic context. This result held for noun-noun ambiguities, but not noun-verb ambiguities, where multiple access occurred in all conditions (including syntactic context). These results are summarized in Table 1.

Table 1. Summary of Results of STL.B's Experiments

Context Type	Ambiguity Type	Outcome
Neutral	Noun-Noun	Multiple Access
Syntactic	Noun-Verb	Multiple Access
Pragmatic	Noun-Noun	Multiple Access
Priming	Noun-Verb	Multiple Access
Priming	Noun-Noun	Selective Access

The apparent anomaly lies in the selective access result in this one condition. STLB attribute the result to intralexical priming by the strong associate preceding the ambiguous word (and the organization of the lexicon: see below). It should be noted that the only meaning of "intralexical" in this context that makes sense is actually "intrasemantic": A single *meaning* of the word, and not the lexical representation of the word itself, is primed. So, they assume the appropriate meaning of the ambiguous word is primed by the associated word's meaning and blocks or inhibits the alternate reading.

STLB conclude from this that the results support a modular, autonomous account of the lexical access process. The only contextual effect, selective access of noun-noun ambiguities, was due to intralexical priming, which is local to the lexicon in their view. Second, the results indicate that there are at least two classes of context which interact with word recognition in different ways. Third, the difference in the results for noun-noun and noun-verb ambiguities suggest that syntactic information is encoded in the mental lexicon. This point is obvious, but the question is *how* syntactic information is encoded. It is possible that a word's syntactic class is encoded with the lexical representation or with the meaning representation. The distinction will become clear when we see their model, which chooses the former, and ours, which chooses an intermediary position. Finally, the results suggest that studies which illuminate the time course of comprehension processes are essential to decoding the structure of the processor(s).

STLB's MODEL

STLB's model is a combination of Morton's (1969) logogen model and Collins and Loftus' (1975) spreading activation model. A lexical logogen governs recognition, and is connected to semantic memory where it activates its meaning(s) via spreading activation. The meaning nodes are accessed along pathways from the lexical nodes in the order of relative activation levels. The meaning nodes may be primed by the access of words highly related to one meaning, which is the only exception to the automaticity and autonomy of lexical access. They posit that if there are large differences in activation due to frequency or priming, then selective access obtains.

In order to account for the difference in noun-noun vs. noun-verb results for semantic context, they posit that nouns and verbs have different nodes with identical recognition procedures in the lexical network (See Figure 1). Now, the story goes, for noun-verb ambiguities with one meaning primed, both nodes get recognized because they share all the same features, and both meanings are accessed. In the noun-noun case, if one meaning is primed, that *pathway* is followed first. Note that this explanation implies serial evaluation of the possibilities in the noun-noun priming case.




Figure 1: STLB's model of lexical access.

A CONNECTIONIST MODEL OF LEXICAL ACCESS

First, a short introduction to connectionism. Connectionist models consist of simple processing units connected by links. The units have a small set of *states* (not used in the following model), a bounded *potential* (we use the range 0 to 1), an *output*, which for our purposes is just a thresholded potential*, a vector of *inputs*, and functions for computing a new state, potential and output from the old ones and the inputs. There are no constraints on the functions that can be used, though they should be kept simple. (It is an important research topic at the moment to discover what constraints on the functions can be reasonably assumed without losing computational ability.) The basic idea is that a unit stands for a value (the infamous "grandmother cell"), and collects inputs from other units which represent evidence for that value, positive or negative. The links between the units are weighted at the input sites, reflecting the importance to the receiving unit of the evidence from that link. Thus much of the information is contained in the connections between units (hence the name "connectionism"). A unit's output represents its confidence in the hypothesis that its value is represented by its input. Thus the typical way to go about building connectionist models is to first decide on what the elements of the domain are that we want to model, choose a way to encode those as units, and then to wire the units together in such a way as to encode *constraints* between the elements. Finally, we must choose an appropriate function for combining the evidence.

Our model for the lexical access process is shown in Figure 2. We show the network for the word "deck", since it is at least four ways ambiguous, with two noun meanings and two verb meanings. The network for a noun-noun ambiguous word would just consist of the left half of this network, (right half for verb-verb), and a noun-verb ambiguous word would just have the outer "V" of seven nodes. The lowest node represents the lexical item and is assumed to be activated by a phoneme or letter recognition network (such as the one described in McClelland & Rumelhart, 1981). The top row of nodes represent the various meanings of the lexical item and are assumed to be connected into a sentence

*Official versions of the theory require that there be only integer outputs from 0 to 10, in order to model the small number of bits that can be encoded in neuronal firing frequency. We are not purists in this respect.

processing and/or an active semantic network. The lexical node activates its meaning nodes through a discrimination network, starting with the grossest distinctions possible, then progressively finer ones. Note that the most efficient way to do this is to make two-way splits between large classes of alternatives (divide and conquer), if possible (but we don't assume all splits are two-way), since the inhibitory connections are minimized this way*. We assume that syntactic information is more discriminatory than semantic information, i.e., that the distinction into "noun" and "verb" divide the possibilities up more than divisions based on meaning.

The alternatives at any discrimination inhibit one another, so that one path through the network eventually "wins" and the meaning nodes that the other paths support fade away. This is the decision process. We assume that this process is driven by feedback to the meaning nodes from higher levels in the network. In the case of a biasing sentence, this would be from higher level nodes representing the role that meaning could play in the sentence, as in, for example, the Cottrell & Small (1983) model of sentence processing. (We also assume there is not a direct link to such role nodes.) In the case of semantic priming, we assume the meaning node is directly primed by a node representing the relation of the priming meaning to this meaning, as in the Collins & Loftus (1975) model of semantic priming. The unfortunate meaning node that does not get top down feedback (or does not get as much) will not be able to provide as much feedback to the pathway nodes which activated it, and its pathway will be inhibited by the pathway nodes that do get more feedback.

In order to account for the modular nature of lexical access, we had to make two simple assumptions about the units. We assume that the units are thresholded (i.e., they can collect activation but they will not fire until they cross threshold, as in Morton's (1969) "logogen" model) and that top down links have lower weights than bottom up links. A unit may thus be activated above threshold by bottom up evidence, but not by top down evidence. This combination of threshold and weighting acts as a barrier to top down information affecting lower level processes by itself, such as recognition. It may come in to play, however, *after* recognition of the lexical item has begun, in the decision process. This assumption is independently motivated at all levels of our networks by the need to prevent top down activation from hallucinating inputs.

*For n nodes to be mutually inhibitory, we need $(n-1)*n$ inhibitory connections. If we arrange them in a binary tree, we need $n-1$ inhibitory connections, but $2*n-1$ units, so we are making a connection/unit tradeoff. This is motivated by the observation that we can't assume the network is pre-wired (in humans). The connectionist model of forming connections involves recruiting units that are on the path between two units (Feldman, 1982). Thus, by conserving connections, we are really conserving units as well.

An interesting feature of this network is that the meanings themselves are not mutually inhibitory. When one considers constraints between units, there is no functional reason to assume that a particular *meaning* in isolation from its source (a particular lexical item) is not compatible with another meaning. However, it is reasonable to assume that the *assignments* of different meanings to the same use of a word is inconsistent. Indeed, if the meanings themselves were mutually inhibitory, we would expect that a word with the same meaning as an inappropriate reading of a previous word in the sentence (assuming the meaning node is shared) would be harder to process than a control word. For example, this would imply that it should be hard to understand "I had a *ball* at the *formal dance*". Our model would predict, however, that people would be slower at processing sentences such as "I had a *ball* at the *ball*".

AN EXAMPLE RUN

We present the result of running the model using the ISCON simulator (Small *et al.*, 1982) in Figure 3. It will be helpful to refer to Figure 2 to understand the trace. We include a driver node, m1 (not shown), that provides constant feedback to SHIP-FLOOR throughout the simulation. (In a complete model this would be a node representing one of the types of SHIP-FLOOR. For example, m1 could be PART-OF-SHIP, activated by the context prime "ship's"). The units average their input from three sites, bottom up, top down, and inhibitory. The first two sites take the maximum of their inputs, and the inhibitory site uses a parameterized arctangent function to enhance the difference in inhibition between two units that are close to each other in activation level. This helps avoid the problem of two units getting into equilibrium without one suppressing the other below threshold. Bottom up weights are 1.0, top down are .5, and inhibitory weights are -0.5. The threshold is set at 0.3. The potential function is similar to the one used by McClelland & Rumelhart (1981).

At step 5, SHIP-FLOOR has been primed by the context prime m1. Now we activate "deck", and continue feeding it for 30 steps. We skip along to step 13 where the semantic discrimination nodes (the "as Xmeaning" nodes) have just fired (not visible at Figure 3's resolution), but their activation has not spread to the meaning nodes yet. Notice that SHIP-FLOOR has been primed now to near threshold. Thus the bottom up activation from "as Nmeaning1" causes it to fire in step 14, while the other meaning nodes have to accumulate more activation for several steps before they will fire. This gives SHIP-FLOOR a chance to increase the relative activation of nodes that are on its feedback path, before the other meaning nodes fire. This allows the nodes on that path to begin to win over their competition so that by step 24, "as Nmeaning2" has been suppressed. This results in CARD-DECK fading from lack of support.

Also, "as Nmeaning1" is no longer inhibited by "as Nmeaning2", so it rises, giving more support to "asNOUN", which then suppresses "asVERB". Later, KNOCK-DOWN and DECORATE fade due to lack of support from "asVERB".

DISCUSSION AND CONCLUSION

This model makes several claims about lexical access. First, decisions within a syntactic class happen "nearer" the meaning nodes than decisions between classes, so the incorrect meaning nodes fade faster when within the same class as its competitors than when its competitors

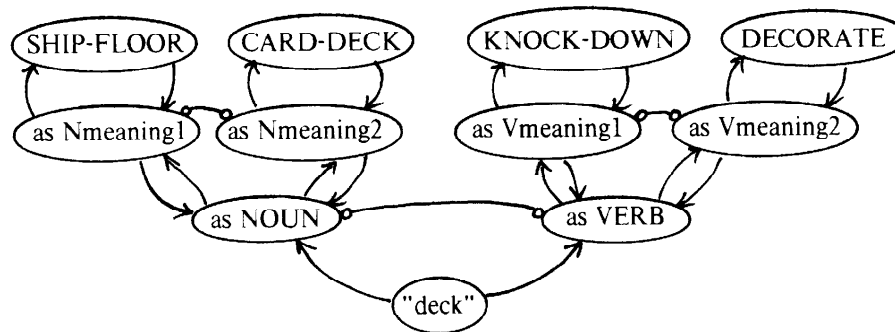


Figure 2: Our model of lexical access.




Figure 3: Trace of the simulation of the network in Figure 2 (x means firing).

are in different classes. Thus noun-noun decisions are faster than noun-verb decisions, as was seen in the sample run. Thus it predicts that verb-verb ambiguities, which have not been tested (to our knowledge) in the psycholinguistic literature, will act like noun-noun ambiguities. However, the STLB study used homonyms (words with unrelated meanings). Verbs tend to polysemy (related meanings). Because this may affect the results, we restrict our claim to verb-verb homonyms.

In order to explain different context effects we have to mention some claims about context. We saw how in our model feedback does not flow freely downward from the priming node (m1) through the meaning node (SHIP-FLOOR) because it is blocked by SHIP-FLOOR's threshold. However, when activation comes up from "deck" through the other nodes, the barrier is broken, and feedback flows down. If we assume that higher levels of processing act the same way, then in the case of pragmatic context, no feedback to meaning nodes would occur before the meaning node actually fired because it is too far away in the network. By this time, multiple access has occurred, and a target word to be named (say, "spade") can take advantage of the priming from all of "deck"'s meanings.

The case illustrated in the sample run was one of priming context with a noun-noun ambiguity (ship's->deck). Here, the contextual priming word is so closely related to one of the ambiguous word's meanings that they are not far away in the semantic network and direct priming of the meaning occurs (eg., "ship's"->SHIP-PART->SHIP-FLOOR). A decision will be reached much more quickly than in the case of pragmatic context, where the feedback has to come from "farther away" (semantically) in the network. Therefore, the model claims that there will be faster decisions in strongly priming contexts. Yet, contrary to STLB, multiple access did occur in our version of a semantic context. We rely on our prediction of the relative speed of ambiguity resolution in different contexts to resolve this. Naming presumably requires at least two stages, recognition and production. The word to be named is presented at the end of the contextually primed ambiguous word. If the decision for the ambiguous word is over before the recognition stage of naming completes, the naming process could not make any use of priming from the alternate meaning of the ambiguous word*. Thus we claim multiple access always occurs, and if the word to be named were presented slightly *before* the end of the ambiguous word, we would see multiple access.

*This claim can be relaxed if we assume our barrier (threshold) is "leaky", that is, with enough top down activation, the meaning node might actually cross threshold before it got bottom up activation. It would then be able to prime the semantic decision node below it to the point where the alternate meaning never gets active. This can be made to happen by using more priming from m1. Our model is therefore in the "chameleon" class with respect to this particular issue.

Finally, in the case of four way ambiguous words such as "deck", we claim that we would see the pattern of results seen in our sample run: In a semantic context, the alternate meaning within the same class would be deactivated first, then the meanings in the other class.

In conclusion, we have designed and built a model of lexical access within the connectionist framework that accounts for the data and makes empirically verifiable claims. This model has several advantages over STLB's in that (1) we don't have to posit nodes with identical recognition procedures, (2) the decision process is motivated by the discrimination network and the difference between nouns and verbs "falls out" of that representation, and (3) it is a computational model. With respect to Artificial Intelligence, we have a parallel model which tackles the major problem of the decision process between the possibly many meanings of a word. An interesting goal now is to design the levels above this which drive the decision process. We think this is a strong case for continued research in the area of connectionist models.

ACKNOWLEDGEMENTS

I would like to thank Michael Tanenhaus and James Allen for helpful comments on this paper. Any errors that remain are mine.

REFERENCES

- Addanki, Sanjaya. A connectionist approach to motor control. Ph.D. thesis, Computer Science Dept., U. of Rochester, 1983.
- Collins, A. M., and Loftus, E. F. A spreading activation theory of semantic processing. *Psychological Review*, 1975, 82, 407-428.
- Conrad, C. Context effects in sentence comprehension: A study of the subjective lexicon. *Memory and Cognition*, 1974, 2, 130-138.
- Cottrell, G. W. and Small S. A connectionist scheme for modelling word sense disambiguation. *Cognition and Brain Theory*, 1983, 6, 89-120.
- Fahlman, S. A. The Hashnet Interconnection Scheme. Technical Report, Computer Science Department, Carnegie-Mellon University, June 1980.
- Fahlman, S. A., G. E. Hinton, and T. J. Sejnowski. Massively parallel architectures for AI: NETL, Thistle, and Boltzmann machines. In *Proceedings of the National Conference on Artificial Intelligence*, Washington, D.C., August 22-26, 1983.
- Feigenbaum, E. A. and Julian Feldman, eds. *Computers and Thought*. New York: McGraw Hill, 1963.
- Feldman, Jerome A., Dynamic connections in neural networks. *Biological Cybernetics*, 1982, 46, 27-39.
- Feldman, Jerome A. and Dana Ballard. Connectionist Models and their Properties. *Cognitive Science*, 1982, 6 205-254.

- Forster, K. I. Levels of processing and the structure of the language processor. In W.E. Cooper & E.C.T. Walker (Eds.), *Sentence Processing*. Hillsdale, NJ: Erlbaum, 1979.
- Foss, D. and Jenkins, C. Some effects of context on the comprehension of ambiguous sentences. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 577-589.
- Holmes, V.M. Prior context and the perception of lexically ambiguous sentences. *Memory and Cognition*, 1977, 5, 103-110.
- Lackner and Garret. Resolving ambiguity: Effects of biasing context in the unattended ear. *Cognition*, 1972, 1, 359-372.
- Lesser, V.R., and Erman, L. D. A retrospective view of the Hearsay-II architecture. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, 1977.
- Marslen-Wilson W.D. and Tyler, L.K. The temporal structure of spoken language understanding. *Cognition*, 1980, 8, 1-71.
- McClelland, James L. and David E. Rumelhart. An interactive activation model of the effect of context in perception: Part I, An account of basic findings. *Psych. Review*, 88,
- Meyer, D. E., and Schvaneveldt R. W. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 1971, 90,
- Morton, J. Interaction of information in word recognition. *Psychological Review*, 1969, 76,
- Sabbah, Daniel. A connectionist approach to visual recognition. TR 107 and Ph.D. thesis, Computer Science Dept., U. of Rochester, April 1982.
- Seidenberg, M. S., Tanenhaus M., Leiman, J. and Bienkowski, M. Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*,
- Shastri, L. and Feldman, J.A. Semantic Networks and Neural Nets. T.R. 131, Dept. of Computer Science, University of Rochester, May 1984.
- Small, S. L., Shastri L., Brucks M., Kaufman S., Cottrell, G., and Addanki, S. ISCON: An Interactive Simulator For Connectionist Networks. Technical Report 109, Department of Computer Science, University of Rochester, Dec. 1982.
- Swinney, David A. (1979), Lexical Access during Sentence Comprehension: (Re)Consideration of Context Effects. *Journal of Verbal Learning and Verbal Behavior*, 1979, 18, 645-660.
- Swinney, D. A., and Hakes, D.T. Effects of prior context upon lexical access during sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1976, 15, 681-689.
- Swinney, David A., William Onifer, Penny Prather, and Max Hirshkowitz. Semantic facilitation across sensory modalities in the processing of individual words and sentences. *Memory and Cognition*, 1979, 7, 159-165.
- Tanenhaus, M., Leiman, J., and Seidenberg, M. S. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, 1979, 18, 427-440.