This is a pre-publication version.

In R. M. Klein & P. McMullen (Eds.), *Converging methods for understanding reading and dyslexia. Cambridge*, MA: MIT Press. (1999; pp. 373-411)

Interactive Processes in Word Identification: Modeling Context Effects in a Distributed Memory System

Michael E. J. Masson

University of Victoria

Correspondence to: Michael Masson
Department of Psychology
University of Victoria
P.O. Box 3050
Victoria, B.C.
Canada V8W 3P5

phone: (250) 721-7536
fax: (250) 721-8929
e-mail: mmasson@uvic.ca
courier: Room A234, Cornett Building

The development of connectionist, or neural network, models has provided a new and promising way of understanding cognitive functions. Reading processes, particularly those associated with the identification of words, have been the subject of extensive study through computational modeling with connectionist systems. The general goal of this enterprise has been to develop a theory of skilled word recognition processes, including accounts of the development of that skill (Seidenberg & McClelland, 1989) and of neurological damage that leads to disorders of reading ability (Hinton & Shallice, 1991; Plaut, McClelland, Seidenberg, & Patterson, 1996; Plaut & Shallice, 1993). Connectionist models of these phenomena have come to rely on a division of labor between different sources of knowledge to account for successful and patterns of impaired word reading performance (e.g., Plaut et al., 1996). The general theme pursued in this chapter is the issue of how knowledge sources interact to bring about fluent word reading.

The proposition that multiple knowledge sources interact in the process of identifying words has been of fundamental importance in the development of models of reading (e.g., Lesgold & Perfetti, 1981; Massaro, 1979; Masson & Sala, 1978; Plaut & Shallice, 1993; Rumelhart, 1977). Not surprisingly, the nature of this interaction is a source of controversy, for it cuts to the heart of the defining principles of differing views of cognitive architecture such as the debate about whether perceptual input systems are modular, and therefore unaffected by higher level knowledge (e.g., Fodor, 1983; Rueckl & Oden, 1986; Seidenberg & Tanenhaus, 1986), or nonmodular and consequently influenced in a direct way by nonperceptual information (e.g., Rhodes, Parkin, & Tremewan, 1993).

In this chapter, I describe a simple computational model of the knowledge sources that are assumed to interact during the course of performing various word identification tasks. The model constitutes a particular view of how this interaction takes place. It provides an account of performance in various word identification tasks, especially the influence of contextual information on task performance. Context effects are a crucial source of empirical evidence regarding the interaction of knowledge sources and they provide constraints on model development.

The model has been implemented with a distributed representation of knowledge and its processing assumptions are in the connectionist or parallel distributed processing tradition. The distributed representation of knowledge leads to an account of word identification phenomena that is qualitatively different from that of models that assume a localist representation. In a localist representation, each concept is represented by a single processing unit, and an arbitrary number of units can be simultaneously active to any degree, depending on assumptions about how the units interact (e.g., Anderson, 1983; Collins & Loftus, 1975; McClelland & Rumelhart, 1981). In contrast, a distributed representation places natural constraints on the simultaneous activation of multiple concepts. In a distributed representation scheme, knowledge about words is represented as connection weights between processing units. This knowledge actually constitutes a "potential" for various known words to come to mind. For a specific word to come to mind, which can be caused by presentation of its visual form, its unique pattern of activation across a set of processing units must be invoked. I discuss below how a word's pattern of activation is instantiated, but for now the important point is that full activation of a word requires setting all of the processing units to the activation values dictated by that word. Thus, it is not possible for two different words to be fully activated at the same time (see McClelland, 1985, for a suggestion regarding how two simultaneously presented words might be processed in a connectionist model). Multiple words are partially activated, however, if the processing units form a pattern that is similar but not identical to the corresponding patterns of activation for those words. Moreover, a change in any subset of the processing units affects the degree of activation of all words known to the system. As shown below, these characteristics of distributed representation provide for a natural account of context effects such as semantic priming.

A second important aspect of the model is that the influence of different knowledge sources operates in a cascaded fashion (McClelland, 1979), providing for continuous change in the availability or activation of knowledge about specific words. The continuous influence of different sources of knowledge is combined over time to allow the model to generate responses. For example, in the task of reading aloud printed words, the model accumulates knowledge about a word's phonology by computing the correspondence between orthographic and phonological patterns, and also by computing the meaning associated with an orthographic pattern, then computing the phonological pattern associated with that meaning. These two sources of influence on the generation of a phonological code operate simultaneously and have a combined influence over time on the construction of a phonological code (see also Plaut et al., 1996; Seidenberg & McClelland, 1989).

This approach to reading a word aloud resembles models that assume there are independent routes to pronunciation, such as a direct route from orthography to phonology and an indirect route either through semantics or an orthographic lexicon (e.g., Besner, this volume; Besner & Smith, 1992; Coltheart, Curtis, Atkins, & Haller, 1993; Paap & Noel, 1991). These multiple-route models characterize word reading as a race between independent routes that yield either the same response or differing responses (as in the case of exception words such as "have"). If the two routes deliver identical phonological responses, the faster route can drive the response. But in the case of exception words the routes produce opposing responses that must be competitively resolved. In the connectionist model described here, as in other connectionist models of word identification, however, the semantic and orthographic routes to phonology are trained to produce compatible phonological patterns even for exception words. Rather than implementing general pronunciation rules in the orthography to phonology route, connectionist models employ learning rules that capitalize on the statistical structure that characterizes the relation between orthographic and phonological patterns (e.g., Seidenberg & McClelland, 1989). In principle, the correct pronunciation for both regular and exception words can be generated by the orthographic route in connectionist models.

In the next section of this chapter, a full specification is provided of the connectionist model that I have been using to account for how knowledge sources interact to produce a range of context effects. A review is then provided of recent applications of the model to simulate semantic context effects (including tests of informational encapsulation), semantic ambiguity effects, and masked priming effects. Finally, a number of unresolved issues and new directions for development of the model are considered.

Architecture of the Distributed Memory Model

The connectionist model I have developed (henceforth referred to as the distributed memory model) adopts the general architecture proposed by Seidenberg and McClelland (1989), in which separate sets or modules of processing units are devoted to representing a word's orthographic, phonological, and semantic information. Although the original version of the distributed memory model (Masson, 1991) consisted of only two sets of processing units, called perceptual and conceptual, a third set (phonological) was later added to permit simulation of the word naming task (Masson, 1995). The implementation of semantic information in the distributed memory model is constrained so that there is no feedback from semantics to orthography. Thus, there is no top-down influence of semantic knowledge on orthographic processing. Semantic knowledge, once activated by orthography, influences the development of a phonological representation and thereby contributes to performance on word naming tasks. A schematic representation of the model is presented in Figure 1. Flow of activation between modules is indicated by arrows. Furthermore, each unit within the semantic and phonological modules is connected to all other units in its module.

Insert Figure 1 about here

The model is implemented as a Hopfield network (Hopfield, 1982), in which each unit in the network is connected to all other units. The only modification made in creating the distributed memory model is that units responsible for representing orthography are affected only by visual input, not by one another nor by units in the semantic or phonological modules. This constraint on connectivity means that the distributed memory model assumes informational encapsulation of visual word input in the sense proposed by Fodor (1983). Although this restriction on input to orthographic units is not intended as a crucial assumption of the model, it has been useful in the examination of issues associated with context effects, particularly the question of what kind of evidence may be taken as evidence against the informational encapsulation assumption.

Knowledge Representation

A fundamentally important property of the distributed memory model is that word knowledge of all types is represented in a distributed fashion. For example, knowledge about a particular word's orthography is represented as a pattern of activation across all units in the orthographic module. A word's phonology and meaning are similarly represented as patterns of activation across the phonological and semantic units. Each unit can take on one of two values, 1 or –1, so a word's pattern of activation consists of a vector of ±1 values across all the units in the network. No assumptions have been made regarding correspondence of particular processing units or subsets of units to specific letters, letter positions, phonemes, or semantic features. In models such as those of Plaut et al. (1996), Plaut and Shallice (1993), and Seidenberg and McClelland (1989), the objective has been to develop a computation model of the mapping of orthography to phonology or orthography to semantics for specific English words. In contrast to these models, the goal for the distributed memory model has been more modest; namely, to examine general properties of a system in which multiple knowledge sources converge to generate a representation in response to input. For this purpose, it has not been necessary to incorporate details of, for example, English orthography and phonology. Neither has it been necessary to posit a particular set of semantic features (cf., Hinton & Shallice, 1991; McRae, de Sa, & Seidenberg, 1997; Plaut & Shallice, 1993) in representing word meaning. It has been sufficient to use arbitrarily (even randomly) constructed patterns to represent each word's orthography, phonology, and semantics.

The use of arbitrary patterns to represent word knowledge does not imply a rejection of the ideas put forward in models designed to capture aspects of language such as grapheme-phoneme correspondences or the specification of semantic features. To the contrary, the distributed memory model embraces the notion, developed by Seidenberg and McClelland (1989), that pronunciation of both the regular and exception words can, in principle, be computed by a single pathway between a distributed representation of orthography and a distributed representation of phonology. Because the goal of developing the distributed memory model was not to account for aspects of linguistic competence such as these, it has been possible to use a simpler representational scheme involving arbitrary patterns of activation for each word.

Learning

The distributed memory model uses a learning rule derived from Hebb (1949) and discussed by Hopfield (1992) in which the connection weight between any pair of units is changed as a function the activation states the units take on when a pattern to be learned is instantiated. When a specific pattern of activation (representing a particular word) is learned, the connection weight between each pair of units is adjusted. In particular, if the pattern of activation calls for two units to be in the same activation state, the connection weight between them is increased; if the pattern dictates different activation states for the two units, the connection weight between them is decreased. The formula for this learning rule is presented in the Appendix.

The effect of this learning rule is to increase the connection weight between two units that typically are in the same activation state across the various patterns of activation (words) that the system learns. For two units that typically are in different states, the connection weight is decreased (i.e., becomes negative). If, across the entire vocabulary of words, a pair of units is about equally often in different states and in the same state, the net weight change will be close to zero. Thus, the learning rule allows connection weights to capture correlations between pairs of units that hold over the full set of words.

The disadvantage of using this simple learning rule is that the distributed memory model is capable of learning only a small set of words. I have found, however, that even with a vocabulary of 6 to 12 words, the model is able to provide an account of a range of interesting word identification phenomena.

Pattern Completion

A crucial feature of Hopfield networks is their ability to complete a partial pattern instantiated in the network to form a previously learned pattern of activation (Hopfield, 1982). In particular, Hopfield networks are attractor networks in that learned patterns of activation lie at the bottom of basins of attraction in a multidimensional space consisting of all possible patterns of activation across the network's units. The process of pattern completion begins by instantiating part of a previously learned pattern of activation across a subset of the network's units. This subset of units is fixed, or clamped, in this pattern, then the network is then allowed to update the activation states of its other units. A unit's activation state is determined by computing the input coming into that unit from all other units in the network. The input received by a unit will cause it to take an activation state of either +1 or –1. The formula for computing the input to a unit and the rule for using input to determine a unit's activation state are provided in the Appendix. Over time, as units are sampled and their states computed, the network moves, step by step, into the learned pattern dictated by the initial, partial pattern that was instantiated in the network.

Once the network has fully instantiated a learned pattern, it will normally remain in that state until a new partial pattern is instantiated in the network. Thus, learned patterns are sometimes referred to as stable states. The progress of the network as it moves down a basin of attraction, toward a learned pattern, can be tracked by computing a measure known as energy. As the network moves more deeply into a basin, energy is reduced (takes on a larger and larger negative value), achieving a maximal negative value when the bottom of the basin is reached and the learned pattern is fully instantiated. The formula for computing energy is provided in the Appendix. An important interpretation of the energy of the network is that of goodness of fit: a large negative value of energy indicates that the current pattern of activation across the units constitutes a good fit between the current pattern and the connection weights in the network. That is, pairs of units that have positive connections tend to be in the same state and those with negative connections tend to be in opposite states. Energy can be computed across the entire network or within a subset of the modules, depending on which sources of information are assumed to be used in performing a particular task.

Word Identification

The distributed memory model simulates word identification as a pattern completion process in which a word's orthographic pattern is instantiated in the network and the remaining processing units are updated to complete the target word's pattern. In the applications for word identification tasks described here, various criteria have been applied in simulating responses, depending on the nature of the task. For example, simulation of the word naming task involves instantiating a word's orthographic pattern, then updating the semantic and phonological units until the latter set of units forms a pattern of activation that corresponds to the target word's phonology. The number of updating cycles required to simulate a response is taken as a measure of response latency. Although it is not clear what function ought to be used to map number of updating cycles to response latency, it has proven adequate to assume there is a linear relation between number of updates and response latency.

Converging sources of influence. As indicated in Figure 1, the model includes two sources of influence on phonological units that originate outside the phonological module: an influence from orthography and another from semantics. Similarly, the instantiation of a word's meaning is affected by two sources outside the semantic module: orthography and phonology. With two converging sources of input influencing the reconstruction of a phonological or a semantic code, questions arise as to the relative strength and the timing of the two influences.

Two factors together determine the strength of a module's influence on a particular unit in the network (either within that module or within a different module). The first is the number of units in a module. The more units in a module the larger will be that module's contribution to the net input to a unit. In the simulations described below, the typical numbers of units in the three modules were 130 orthographic units, 80 semantic units, and 40 phonological units. The values were determined primarily by practical constraints. First, using a large number of orthographic units is one way of making it likely that the rest of the processing units will form the correct pattern of activation when a particular word's orthographic pattern is instantiated. (A similar effect could be accomplished with fewer orthographic units by implementing a form of gain control on the weights connecting orthographic units with other units.) Second, to keep the time needed for running simulations within reasonable bounds, the number of phonological units was made relatively small. By using fewer units, the average number of processing cycles needed to reach a stable state across the entire phonological module was kept to a reasonable value. Finally, the choice of the number of semantic units was determined in part by the need to provide for different words that have similar semantic patterns so that semantic priming effects could be simulated. With 80 semantic units, there is freedom to examine a range of semantic similarity.

The second factor that determines a module's influence on a unit in the network is the rate at which the module's units are updated. Different modules have the potential to have their units updated at different rates. In the simulations reported by Masson (1995), it was assumed that when a word was visually presented, its orthographic units immediately took on the relevant pattern of activation, then the phonological and semantic units began updating. The immediate instantiation of a word's orthographic pattern was used only for convenience; in later simulations, orthographic units were updated over time rather than immediately shaped to the correct pattern. It is also assumed that phonological units update at a higher rate than semantic units, allowing phonological units to have an early influence on instantiation of the word's meaning. This approach is consistent with empirical results suggesting that phonological recoding of words occurs early during word reading and contributes to the access of meaning (e.g., Lukatela & Turvey, 1994; Perfetti & Bell, 1991; Perfetti, Bell, & Delaney, 1988; Van Orden, Johnston, & Hale, 1988; Ziegler & Jacobs, 1995). Because the orthographic module is larger than the phonological module and is either fully instantiated at the start of a trial or updated at a higher rate than units in any other module, however, orthographic units will have the dominant influence on units in the network. In this sense, the model is sensitive to findings that indicate orthography can have a direct influence on access to word meaning (Daneman, Reingold, & Davidson, 1995; Jared & Seidenberg, 1991).

Lexical access vs. pattern formation. Nondistributed, or localist, models of word identification typically assume that word knowledge is represented in some form of lexicon (Coltheart et al., 1993; Forster, 1976; Jacobs & Grainger, 1992; Johnson & Pugh, 1994; Morton, 1969). The lexicon contains an entry for each known word and word identification consists of accessing the appropriate entry in the lexicon. Localist models provide for the possibility that multiple words can be accessed or activated at the same time. When multiple lexical entries become active simultaneously, one means of determining which active entry corresponds to the target word is to invoke an inhibitory process that permits the correct entry to be selected against the background of similar, active entries. For example, in the interactive activation model, it is assumed that lexical entries mutually inhibit one another so that the entry most strongly supported by orthographic evidence eventually wins out (e.g., Jacobs & Grainger, 1992; McClelland & Rumelhart, 1981).

In the distributed memory model, and in other models that assume a distributed representation of word knowledge, knowledge about multiple words can become "activated" simultaneously, but only partially so. A distributed network can move into patterns of activation that partially match a number of known words. A word is activated to the extent that its pattern is instantiated in the network's units. Only one word at a time, however, can have its full pattern instantiated across the entire network. Thus, a distributed network provides a natural solution to the problem of distinguishing among a number of competing candidates: as the network forms a pattern of activation in response to presentation of a word, it eventually settles into the target word's pattern and in so doing moves away from patterns associated with other words. Cohorts, neighbors, or other related words that might have become partially activated during early processing of a target word lose their activation as the target pattern takes hold. This natural suppression of related words is a consequence of the shared representational space constituted by the network of processing units.

To illustrate this property of the distributed memory model, a simulation was run in which a vocabulary of nine items was learned. The items constituted three triplets, each consisting of a target word and two related words. One related word was similar to the target only with respect to orthography and the other was similar to the target only with respect to meaning. In the distributed memory model, similarity between words is defined as similar patterns of activation across a set of units. Two words that are, for example, orthographically similar, would have a high proportion of orthographic units in the same activation state (± 1) when their respective orthographic patterns are instantiated in the network. In the simulation described here, orthographically similar pairs shared the same activation states in 90 of the 130 orthographic units. Unrelated words shared activation states in an average of 65 orthographic units (half of 130). Similarly, words that were semantically related had the same activation states in 55 of the 80 semantic units. After learning, each target word was presented to the network multiple times. On each presentation, the network was run for 2000 updating cycles, allowing the network to approach formation of the full pattern of activation for the target word.

To simulate the construction of a word's orthographic pattern over time, the orthographic units were not clamped to the appropriate pattern at the beginning of a trial. Instead, they were initially set to a random pattern, just as all other units in the network, then randomly sampled for updating with high probability (0.60). When sampled, an orthographic unit was clamped to its appropriate state. If the orthographic module was not sampled on a cycle, then a unit from the phonological module was sampled with probability 0.60, otherwise a semantic unit was sampled. Thus the simple probabilities of sampling orthographic, phonological, and semantic units were 0.60, 0.24, and 0.16, respectively.

The degree of activation over time (updating cycles) of orthographic, phonological, and semantic aspects of a target word is shown in Figure 2a. Activation of knowledge about a word was taken as the proportion of units in a module (e.g., orthographic) that were in states that correspond to that word's learned pattern. Once a word is fully instantiated in the network, all units meet this criterion and so the activation value reaches its asymptote of 1.0. The effect of using different probabilities for sampling units from the three modules is apparent in Figure 2a. The relatively high sampling rate for the orthographic units resulted in a rapid rise to asymptotic activation for those units. The small number of phonological units combined with a higher sampling rate than the rate for semantic units, produced a rapid rise in activation of a target's phonological representation. Both the orthographic and phonological patterns were fully formed before the semantic pattern was completely instantiated.

Insert Figure 2 about here

The activations of each type of knowledge (orthographic, phonological, and semantic) is shown separately in Figures 2b-2d. Each plot shows activation for the target word and its related words. Only external input affects the orthographic units, so only the target and its orthographically related word show any sign of activation. The related item's orthographic pattern reaches a level of activation that corresponds to its degree of similarity to the target's orthographic pattern. In the phonological module, the orthographically similar word's phonological pattern is initially activated as the orthographic units begin to form the target's pattern. This activation occurs because the orthographically similar word's orthographic pattern is partially instantiated and the target's pattern is not yet completely instantiated. As the target's orthography becomes more fully established, however, the phonological module is dominated by the target, and activation of the orthographically related word's phonology dies away. No attempt was made to incorporate correlations between orthographic and phonological patterns in the items learned by the network, so the target's orthographic pattern offers no support whatever to the phonological pattern of its orthographic neighbor.

In the semantic module, the semantically related word's pattern of activation grows along with the target's activation, although it reaches a lower asymptote. The semantic pattern of the orthographically related word is also activated during the early stages of processing the target, although this activation eventually dies out. Although the orthographic neighbor has no stronger semantic similarity to the target than does an unrelated word, the pattern of activation in the orthographic units (which are somewhat compatible with the orthographic neighbor) initially push the semantic units toward a pattern that is similar to the orthographic neighbor. As the target's pattern comes to dominate the network, however, any semantic activation of semantically unrelated words dissipates.

In general, the activation of the relevant aspect of a related word grows to a moderate asymptotic value, but does not subsequently drop to the original resting level. Related words will be simultaneously active with the target just to the extent that they share patterns of activation with the target. Although activation of relevant aspects of related words does not die away, Figures 2c and 2d make it clear that temporary activation of other aspects of related words can be short-lived and does eventually fall back to resting level as the target's pattern comes to dominate the network.

This elimination of activation occurs despite the fact that the distributed memory model has no mechanism for directly inhibiting activation of word knowledge. Instead, the loss of activation occurs as a natural consequence of the network's distributed architecture. As one pattern comes to dominate the network, unrelated patterns recede gracefully into the background.

Applications of the Distributed Memory Model

In this section I provide an account of semantic priming within the distributed memory model, then briefly review a number of applications of the model related to priming and other semantic effects in word identification. Then I describe a new set of simulations in which the model was used to account for results from masked priming experiments.

Semantic Priming

Semantic priming effects generally involve enhanced performance on a target word (e.g., shorter response latency) when presentation of the target is preceded by presentation of a related prime word relative to when the prime word is unrelated to the target (e.g., Meyer, Schvaneveldt, & Ruddy, 1972; Neely, 1977). In the distributed memory model, semantic priming effects are a natural consequence of the assumption that semantically related words consist of similar patterns of activation across the semantic units. Priming arises because presentation of a prime word causes the semantic units to begin to form that word's pattern of activation. When a related target is then presented, the pattern of activation in the semantic units is closer to the target word's pattern than would be the case had an unrelated prime, or no word at all, been presented. Therefore, fewer changes need be made to the semantic units to form the target word's semantic pattern (see Sharkey 1989, 1990, for a similar account of semantic priming). Even if responding to the target does not require the target word's semantic pattern to be fully formed, the head start into the target's semantic pattern will be beneficial for various measures of the model's progress toward identifying the target (e.g., settling of the phonological units to simulate a word naming response).

Semantic vs. associative priming. Before discussing the distributed memory model's account of specific experimental results involving semantic priming, some consideration of the distinction between semantic and associative priming is warranted. Semantic priming refers to cases in which words that have similar meanings prime one another in a word identification task. In contrast, associative priming refers to priming effects that obtain when words are associatively related, as measured by free association norms (e.g., Postman & Keppel, 1970), but not semantically related (e.g., milk and cow). Shelton and Martin (1992) found that under conditions intended to prevent the use of expectancy and post-lexical checking, associatively related word pairs generated a priming effect in lexical decision, whereas semantically related pairs did not.

The Shelton and Martin (1992) result poses a problem for the class of models, including the distributed memory model, that assume semantic priming effects are a consequence of similarity in the patterns of activation for semantically related words. More recent empirical studies, however, have produced evidence for automatic semantic priming. First, Moss, Ostrin, Tyler, and Marslen-Wilson (1995) found that under conditions like those used by Shelton and Martin, semantically related words that share a functional relationship (e.g., broom-floor), but are not associatively related, produce a reliable priming effect. Thus, automatic semantic priming can occur in the absence of an associative relationship. Second, McRae and Boisvert (1998) showed that pairs of words that were rated by subjects as highly similar in meaning, but that were not associatively related, produced reliable semantic priming using the Shelton and Martin procedure (see Thompson-Schill, Kurtz, & Gabrieli, 1998, for a similar result). McRae and Boisvert also replicated the lack of a priming effect with the items originally used by Shelton and Martin and showed that those pairs were rated as less similar than the pairs that yielded a priming effect. These new results support the proposal that semantic priming effects arise from semantic similarity between pairs of words and therefore support accounts of priming effects that are based on semantic similarity.

Influence of an intervening item. An important implication of the distributed memory model's representation of word knowledge is that priming effects occur because processing of the prime moves the network's pattern of activation into a state that is favorable to the upcoming target. This basis for priming is susceptible to interference from presentation of an event that intervenes between the prime and the target. In particular, if a word were to be presented just after a prime, but before a related target, and if that intervening word were unrelated to either prime or target, then processing of that word would move the network's pattern of activation away from the pattern that was established by the prime. That is, the preliminary work done by the prime would be dismantled in the course of processing the intervening word, thereby reducing or eliminating the priming effect. Whether the effect is entirely eliminated would depend on the amounts of processing time devoted to the prime and the intervening word. As more time is devoted to the intervening word, the effect of the prime begins to disappear.

The distributed memory model is not the only type of model that predicts this effect of an intervening word on semantic priming. This phenomenon is also predicted by the compound cue model proposed by Ratcliff and McKoon (1988, 1995). In that model, binary decision tasks, such as lexical decision, are performed by assessing the familiarity in memory of a target stimulus. In paradigms that involve presentation of a prime in conjunction with a target item, the prime and target are assumed to form a compound stimulus that becomes the object of the familiarity assessment. Semantic priming effects in this model emerge from the relatively high degree of familiarity of a prime-target compound that consists of a related pair of words, compared to the familiarity of a compound consisting of two unrelated words. By presenting an unrelated word between a related prime-target pair (e.g., cat-spy-DOG), a new compound is formed consisting of the unrelated, intervening item and the target. The familiarity of this compound would be about the same as for any pair of unrelated words, leading to the disappearance of the priming effect.

Empirical evidence regarding the effect of an intervening word has been mixed, with some studies finding that priming still occurs (e.g., McNamara, 1992; Meyer et al., 1972) and others finding that priming is eliminated (e.g., Dannenbring & Briand, 1982; Ratcliff & McKoon, 1988, 1995). Using a word naming task, I found that an unrelated intervening stimulus reduces or eliminates semantic priming and reported a simulation of the result using the distributed memory model (Masson, 1995). These effects are not simply due to the presentation of any stimulus between the prime and target because robust priming effects occur if the intervening item is a neutral stimulus such as a row of <u>x</u>'s or a word that occurs on many trials (e.g., <u>ready</u>).

The intervening stimulus effect is important not only because it supports the predictions of the distributed memory and compound cue models, but also because it is contrary to what would be expected by the Collins and Loftus (1975) model of semantic priming that was based on localist representation of knowledge. In that model, activation spreads automatically from a prime word to all related words, independently of activity in other parts of the network. Thus, presentation of an unrelated intervening word should have no effect on activation that is spreading in a distant area of the network.

A more sophisticated localist representation model, the ACT* model of Anderson (1983), can account for the disruptive effects of an intervening item in the following way. Activation of a prime is assumed to decay once the prime stimulus disappears, unless attention to the prime is maintained. The pattern of priming effects found when the intervening stimulus is an unrelated word versus a neutral stimulus can be explained if it is assumed that activation of the prime decays during processing of an unrelated intervening word, but is maintained by attention when the intervening item is a neutral stimulus. A similar proposal involving automatic spreading activation combined with attentional allocation in semantic memory was made by Posner and Snyder (1975). Thus, although the intervening stimulus effect does not provide a means of discriminating between the distributed memory and compound cue models on one hand, and all versions of spreading activation models on the other hand, it does constitute support for important predictions of the first two models while contradicting predictions of classic spreading activation models.

A related result that might be taken as a challenge to the distributed memory model was reported by Joordens and Besner (1992). In a continuous word naming task, they embedded in the list pairs of target words that were semantically related but separated by an intervening unrelated target. All targets were named as quickly as possible, and a small but reliable priming effect was observed despite the fact that an unrelated word intervened between related targets. This result was problematic for the original version of the distributed memory model (Masson, 1991) because that model simulated only a generic word identification task and did so by requiring the entire set of conceptual units to settle into a stable state. Any advantage created by a related target would be completely destroyed by processing and responding to an unrelated word because the latter word's pattern of activation would completely take over the conceptual units, leaving no vestige of the related target's processing.

A more realistic simulation of word naming was developed in the modified version of the distributed memory model (Masson, 1995), as described above. In the modified version of the model, settling of phonological units is the criterion for word naming. Because the phonological units are fewer in number and update at a faster rate than conceptual units, naming an intervening word did not completely change the pattern of activation in semantic units created by the previous target. Therefore, when the critical target was presented, some of the work in the semantic units done by the earlier related target was still in place. Processing of the related target was able to take advantage of this weak semantic pattern, producing a small priming effect.

Evidence for top-down effects of priming. A fundamental question that models of priming must consider is whether or not conceptual knowledge has a direct influence on lower levels of processing (e.g., development of an orthographic representation or identification of letters). On the view of interactive models, such as the interactive activation models of McClelland and Rumelhart (1981) and McClelland (1991), such top-down influences are a key part of the models' architecture. For example, in the interactive activation model, McClelland and Rumelhart accounted for superior identification of a target letter when embedded in a word as compared to a nonword by assuming that activation is passed from the level of word units back down to the level of letter units. Others have proposed that top-down activation of this kind does not occur, and that higher level knowledge interacts with perceptual information only during a decision stage, as in the fuzzy logical model of perception (e.g., Massaro, 1979, 1989; Rueckl & Oden, 1986). The latter perspective is compatible with the notion of informational encapsulation (e.g., Fodor, 1983), in which it is assumed that input modules are affected only by direct perceptual input, but not by higher level knowledge.

The distributed memory model, as described here and in Masson (1991, 1995) adopts the view that conceptual knowledge does not affect perceptual processing. As shown in Figure 1, the model's architecture does not provide for any influence of semantic or phonological knowledge on the development of an orthographic pattern. Only direct visual input affects the state of the orthographic units. Although this aspect of the model could easily be modified, it has been instructive to examine the model's performance under the assumption that no top-down influences are in effect. This version of the model has been particularly useful in the evaluation of two results that could be taken as support for the proposition that higher level knowledge influences perceptual processing.

First, it has been found that context effects on lexical decision and word naming tasks are amplified when targets are degraded (e.g., Becker & Killion, 1977; Besner & Smith, 1992; Borowsky & Besner, 1991; Meyer, Schvaneveldt, & Ruddy, 1975). To examine the model's performance under conditions of degraded input, a version of the model like that used to produce the data in Figure 2 was used to simulate four word naming conditions representing a factorial combination of related versus unrelated prime and clear versus degraded target. Target degradation was simulated by probabilistically setting an orthographic unit to its correct state when it was sampled for updating. The result of this perturbation in the development of a target's orthographic pattern was generally to slow the rate at which that pattern was formed, and therefore, to increase the number of cycles to settle the phonological pattern, just as degradation generally slows responding in humans. The model also produced an enhanced effect of context when the targets were degraded, as shown in Figure 3. This interaction follows the pattern observed in human data and is generated because the noisy orthographic influence on phonology gives the semantic units a greater opportunity to influence the state of the phonological units. Thus, the effect of semantic priming is amplified under conditions of degraded visual input.

Insert Figure 3 about here

It should be acknowledged that the interaction between context and degradation may not be as straightforward as depicted here. Recent evidence reported by Stolz and Neely (1995) indicates that there are conditions under which context and degradation do not interact in the lexical decision task. In particular, these two factors were additive when the proportion of trials on which the prime was related to the target was low (.25). The potential for degradation and context to be additive poses a significant challenge to connectionist models in general because of their highly interactive character. A possible solution would be to introduce a mechanism that would, under certain circumstances, delay inputs from orthographic to other modules until the orthographic pattern of activation reaches some criterion level of completeness (e.g., Joordens, Masson, & Besner, 1995). This approach would create a stage-like process that might result in additivity between degradation and context.

A second result that supports the proposal that context influences perceptual processes in a top-down manner is the demonstration that accuracy in identifying words, measured using signal detection, can be increased by semantic priming. Following Farah's (1989) argument that signal detection measures of sensitivity reflect perceptual processing, whereas measures of bias reflect attentional and semantic processes, Rhodes et al. (1993) showed that word/nonword discrimination accuracy, measured using A', the nonparametric signal detection version of sensitivity, was significantly greater when the target was preceded by a related prime. Rhodes et al. interpreted this result to mean that priming semantic knowledge had improved the perceptual processing of the targets, constituting a clear example of a top-down influence of context.

Contrary to that conclusion, however, Norris (1995) was able to simulate the effect of context on sensitivity using a criterion-bias model in which each word is represented as a logogen and the effect of a related context word is to reduce the recognition threshold for logogens related to that context word. In the criterion-bias model, context has no effect on the acquisition of stimulus information but the lowered threshold means that less stimulus information is required for a primed logogen to reach the level of activation needed to exceed that threshold. Crucially, a prime lowers the thresholds only of words related to the prime, so in the case of an unrelated prime the target word's threshold is not lowered. The resulting difference in the amount of stimulus information needed to reach threshold in the related and unrelated conditions produces the observed sensitivity effect.

Like the criterion-bias model, the distributed memory model is also able to simulate the sensitivity effect in word identification. Masson and Borowsky (1998) simulated the lexical decision task used by Rhodes et al. (1993) with the distributed memory model by assuming that decisions in this task are based on a feeling of familiarity generated when the target is processed (e.g., Balota & Chumbley, 1984; Besner & Johnston, 1989). We assumed that familiarity corresponds to how well the network's current pattern of activation across all processing units conforms to the constraints imposed by connection the weights. As described above, the energy function provides a measure of this constraint satisfaction (see the formula in the Appendix) and we assume that larger negative energy values correspond to a feeling of familiarity.

Using familiarity (energy) as a basis for the model's lexical decisions, we found that the model produced a robust effect of semantic priming on sensitivity in the simulation of the masked lexical decision task. The reason for the model's prediction is that familiarity was computed across all units and connections in the network, thereby integrating all sources of knowledge in the model: semantic, phonological, and orthographic. Without ever directly affecting perceptual processes or the formation of the orthographic pattern of activation, semantic knowledge had a strong influence on lexical decisions because that knowledge converged with other knowledge sources, including orthographic knowledge, to determine the value of energy. This approach is compatible with the fuzzy logical model of perception (Massaro, 1979; Rueckl & Oden, 1986) that assumes semantic and perceptual information are integrated at a decision stage, without semantic knowledge affecting perceptual processing.

Masson and Borowsky (1998) also conducted two experiments using a probe-matching identification task. In this task, target words were presented briefly and followed by a mask, then a probe word. The task was to decide whether the probe matched the target. The targets were preceded by a related or unrelated prime in the form of either a word or a line drawing. Both types of related prime, word or line drawing, led to increased accuracy, as measured by signal detection. The distributed memory model was used to simulate this effect by assuming that degraded information about the target is held in a temporary buffer and compared against the fully instantiated representation of the probe word. If the similarity between the buffer contents and the probe reached some criterion, a positive decision was made.

The effect of semantic priming in this case was to allow the model to develop more of the target's pattern of activation in the semantic units during the brief time the target was processed. Improved accuracy in the semantic units would also contribute to an improvement in the phonological units. When the target's pattern was loaded into the buffer at the completion of target processing, the pattern would be more accurate with respect to the target's semantic and phonological patterns following a related than following an unrelated prime. When these two sources of information are combined with orthographic information about the prime to make a comparison against the probe word, accuracy is improved because of the enhanced semantic and phonological aspects of the target's representation. This improvement accrues in the model despite the fact that orthographic information about the target is unaffected by priming. Thus, the probe-matching task yielded a reliable effect of priming on accuracy. Once again, however, this result cannot be taken as clear evidence that semantic knowledge affects perceptual processes because the distributed memory model successfully accounts for the result without assuming any top-down influences on perception.

The implications of the model's successful simulation of the sensitivity effect are clear. The results obtained by Rhodes et al. (1993) cannot be taken unequivocally as evidence that semantic knowledge affects perceptual processing. The model provides an existence proof that sensitivity effects can be generated by a modular system, as long as perceptual and semantic sources of knowledge are integrated when making a response.

Lexical Ambiguity

An interesting problem faced by models of word identification, particularly models that use a distributed representation for semantic knowledge, is the resolution of lexical ambiguity. Many words have multiple, unrelated meanings (e.g., bat, chest), yet we typically are able to retrieve one meaning or another without conjuring up bizarre mutations consisting of bits of meaning drawn from the different interpretations of the ambiguous word. The difficulty confronting models of word identification is that a single orthographic input pattern is associated, on different occasions during learning, with different semantic patterns. In the distributed memory model, this difficulty is evident because the model forms basins of attraction that represent a blend of the two meanings of an ambiguous word (Joordens & Besner, 1994) and often fails to form basins of attraction that conform to the learned meanings of that word. When presented with the orthographic pattern of an ambiguous word, the model's semantic units move into a blend state rather than into a state corresponding to one of the word's assigned meanings.

To make it more likely that the distributed memory model would form basins of attraction for the learned meanings of ambiguous words, Borowsky and Masson (1996) implemented a version of the model that had a substantial number of semantic units (140) relative to orthographic units (70). With this version of the model, we were able to simulate results of three experiments we conducted involving ambiguous and unambiguous words in the lexical decision and word naming tasks, as well as a result involving gaze duration during reading comprehension (e.g., Duffy, Morris, & Rayner, 1988; Rayner & Frazier, 1989). The relevant empirical results and the simulation of those results are shown in Table 1.

Insert Table 1 about here

For the word naming task, we found no evidence when testing human subjects of a difference between ambiguous and unambiguous words. In contrast, an advantage for ambiguous words was found in the lexical decision task, but only when pronounceable nonwords were used; no effect was found when nonwords were consonant strings. Gaze duration data from Duffy et al. (1988) and Rayner and Frazier (1989) show that subjects spend more time fixating ambiguous than unambiguous words when these items are presented in neutral context and the two meanings of each ambiguous word are of about equal frequency.

Examination of the model's settling of phonological and semantic units during processing of ambiguous and unambiguous targets showed that the phonological units settled at about the same rate for both types of target (hence there was no effect of ambiguity on word naming). On the other hand, activation of word meaning was initially higher for one of the two meanings of an ambiguous target than for the meaning of an unambiguous target. By 250 cycles of processing, however, this situation had reversed and the meaning of an unambiguous target was more active than either meaning of an ambiguous target. The early advantage of ambiguity in activating meaning arose because on each trial the network was started in a random state and it is more likely that the starting state will be closer to one of the two meanings of an ambiguous word than to the one and only meaning of an unambiguous word. That advantage, however, did not benefit the activation of phonological units because the connection weights between semantic and phonological units were more strongly determined by the learning of unambiguous words than by the learning of ambiguous words. This differential influence of semantics on phonology was due to the training regimen in which the meaning of an ambiguous word was presented once whereas the meaning of an unambiguous word was presented twice to ensure that both types of word were equated with respect to frequency of presentation of their respective orthographic patterns. Thus, the semantic units had a somewhat greater influence on phonological units when unambiguous targets were presented, counteracting the greater activation of semantics by ambiguous targets during the early stages of processing.

In the lexical decision task, responses were based on energy computed across the orthographic and semantic units. If energy reached a threshold value before a specified time limit, the target was classified as a word, otherwise it was classified as a nonword. To simulate the experiment with pronounceable nonwords, nonword orthographic patterns were constructed to be similar to learned words, whereas to simulate the results involving consonant string nonwords, the nonword patterns were very different from the learned words. As shown in Table 1, there was an ambiguity advantage, but only when "pronounceable" nonwords were used.

The reason for the ambiguity advantage is the proximity effect discussed above: the random starting state of the meaning units is likely to be somewhat closer to one of the two meanings of an ambiguous word than it is to the one and only meaning of an unambiguous word. When updating of units begins for an ambiguous word, the influence of the orthographic units will push the meaning units toward the meaning favored by the random starting state of the meaning units. In the early stages of processing, then, ambiguous words will make faster gains in energy than unambiguous words, affording a slight advantage, mirroring the effect Borowsky and Masson (1996) obtained when subjects performed lexical decision task. By using nonwords that permit an easier and earlier discrimination (i.e., nonwords with orthographic patterns very different from those of learned words), however, the ambiguity advantage disappears. In this case, words can be discriminated very early during processing, before the meaning units have been driven very far from their starting state and before the ambiguity advantage has a chance to establish itself.

To simulate gaze duration during comprehension, we took the number of updating cycles for the meaning units to reach a stable state as a measure of gaze duration. This assumption is based on the idea that the eyes remain on a word until sufficient processing of meaning has occurred (the immediacy assumption of Just & Carpenter, 1980). As shown in Table 1, the model took longer to reach a stable state in its meaning units when an ambiguous word was presented, indicating a longer average gaze duration for ambiguous words. The ambiguity disadvantage arises from two different meanings associated with an ambiguous orthographic pattern competing with one another for control of the semantic units.

Although the distributed memory model was capable of capturing the pattern of ambiguity effects across three different tasks by adopting plausible definitions of task performance, the ordering of these tasks along the model's measure of processing time (updating cycles) was quite different from the ordering of these tasks with respect to human latency data. Whereas the model's response latency in the lexical decision task was less than in the naming task, which in turn was less than average gaze duration, just the reverse ordering is true in human data. But the ordering across tasks generated by the model is not easy to interpret because the model is intended to simulate only some of the component processes that contribute to response latencies or gaze durations. For example, in the lexical decision task, the model does not simulate time taken to compute familiarity or preparation and execution of a manual response. In the naming task, the model does not simulate the process of translating a phonological representation into an articulatory code nor the time required to initiate articulation.

The important contribution of the distributed memory model to understanding lexical ambiguity effects is the model's characterization of two counteracting forces. First, as originally suggested by Joordens and Besner (1994), a potential advantage for ambiguous words arises in the model due to the proximity of the random starting state of semantic units to one of an ambiguous word's meanings. By chance alone, the random starting pattern in the semantic units is likely to be closer to one of two meanings of an ambiguous word than to the one and only meaning of an unambiguous word. Second, a disadvantage of ambiguity arises from competition between the two different meanings of an ambiguous word with respect to the pattern of activation that is formed across the semantic units. An ambiguous word's orthographic pattern supports two unrelated semantic patterns and competition between them for the same representational space makes retrieval of meaning less efficient. To the degree that word processing tasks are determined primarily by one or the other of these two phenomena, proximity or competition, either an advantage or a disadvantage due to ambiguity may be obtained.

Masked Priming

An important aspect of the distributed memory model, as in most connectionist models of its kind, is the highly interactive nature of the processing units. The interactive character of this class of model means that additive effects of independent variables can pose a significant challenge to these models, as discussed earlier with respect to context and degradation. The distributed memory model has, however, been able to simulate an additive effect involving word frequency and masked primes. In the masked priming paradigm (e.g., Forster & Davis, 1984), a target word is preceded by a briefly presented prime, which itself is preceded by a mask. The prime, then, is masked both before and after its presentation, making its very presence difficult to detect. Forster and Davis compared identity and unrelated primes using high- and low-frequency words. Because the prime was presented in lowercase and the target in uppercase, the identity prime was orthographically, not physically, identical to the target (e.g., avoid-AVOID). Forster and Davis found that when primes were presented very briefly (60 ms), the advantage of an identity prime over an unrelated prime in a lexical decision task was of the same magnitude for high- and for low-frequency target words (see also Bodner & Masson, 1997).

The distributed memory model was used to simulate this result using a variant of the energy measure, scaled harmony, to discriminate between words and nonwords. Generally speaking, scaled harmony is the inverse of energy, scaled to have a maximum value equal to the number of learning trials, and is described in greater detail in the Appendix.

To simulate the masked priming task, the network was trained on a set of unrelated word patterns. High-frequency words were given four learning trials each, and low-frequency words were given three learning trials each. On each lexical decision trial, a prime was presented for 40 updating cycles, then was replaced by the target. In this simulation, the network was updated as in the simulations whose results are shown in Figures 2 and 3. The value of scaled harmony was monitored during target processing and if the criterion value of .15 was reached before 1,500 updating cycles had occurred, the target was classified as a word. The number of cycles required to reach the criterion was taken as a measure of response latency. With this decision rule, very few errors were made (less than 0.5%), so no error data are reported. Nonword targets were constructed by perturbing the orthographic patterns of learned words. In the case of an identity prime, processing the prime was just the same as giving the network a head start at identifying the target item because the network's orthographic representation makes no distinction between upper- and lowercase, the feature that distinguished primes and targets in the masked priming experiments.

Consideration was also given to the nature of the unrelated prime used when high- versus low-frequency target words were presented. In the case of an unrelated prime, the network moves into a basin of attraction that is not compatible with the upcoming target and additional processing cycles are required for the network to settle into the target's basin of attraction, relative to what would be required had the network been in a random state when the target was presented. The network generally requires more processing cycles to move away from the pattern of activation associated with a high-frequency word and into the target's pattern than it does when moving away from a low-frequency word. Moreover, the network moves more quickly into the basin of attraction for high-frequency targets than it does for low-frequency targets (the basic word frequency effect). Consequently, priming effects should be enhanced when high-frequency unrelated primes are paired with low-frequency targets, but reduced when low-frequency unrelated primes are paired with high-frequency targets.

Forster and Davis (1984) did not specify whether unrelated primes were matched to identity primes with respect to word frequency. Therefore, simulations were run with primes matched on frequency and again with frequency of unrelated primes allowed to vary. When unrelated primes varied in frequency with respect to the target items, an interaction between prime and target frequency was found: there was more priming for low-frequency targets. When unrelated primes were matched to the targets on frequency, however, priming and frequency were additive, as in the Forster and Davis data. The results of the simulation in which primes and targets were matched on frequency are shown in Figure 4. The growth of scaled harmony as a function of updating cycles is shown in Figure 4a, beginning with the presentation of the masked prime 40 cycles prior to the onset of the target. Notice, first, that harmony grows much faster for the word patterns, particularly for high-frequency words, than for the nonword patterns. It is these differences in harmony that produces the discrimination between words and nonwords and the word frequency effect.

Insert Figure 4 about here

The harmony functions for nonword targets also show an effect of priming. In contrast, Forster and Davis (1984) found no identity priming effect for nonwords in the lexical decision task and attributed that lack of priming to the hypothesis that identity priming arises from activation of a lexical entry. On the Forster and Davis account, because nonwords have no entry, there is no basis for priming. Alternatively, it is possible that even the benefit of priming apparent in Figure 4a might be part of human experience, but the increment in familiarity due to priming supports an incorrect response (Bodner & Masson, 1997). Thus, the benefit of priming might be canceled by the conflict in response selection that it creates. The distributed memory model does not have a detailed mechanism for simulating response selection and therefore does not provide an account of this process.

The mean number of cycles required for word targets to reach criterion is shown in Figure 4b. Response latency was less for high-frequency than for low-frequency words, and was reliably reduced by presentation of an identity prime, relative to an unrelated prime of the same frequency. Moreover, word frequency and priming were additive, in line with the results obtained by Forster and Davis (1984) and Bodner and Masson (1997). The reason for this additivity stems from the fact that the identity prime provides the network a head start that is uniform across words of varying frequency, akin to shifting the harmony function (as in Figure 4a) to the left. By using unrelated primes of different frequencies, however, the uniformity of the comparison between identity and unrelated primes breaks down and an interactive pattern emerges from the model, as discussed above.

Limitations and Future Directions

The distributed memory model described here suffers from a number of limitations, some of which were revealed by its application to various word identification paradigms. Other limitations, however, are more general and are shared by other models of the connectionist class, as well as by models of other types. I begin by considering some limitations specific to the distributed memory model and some ideas for overcoming those constraints. I conclude by considering an issue of broader concern to formal models of word reading.

Facilitation versus Interference

In a recent critique of the distributed memory model (and attractor network models more generally), Dalrymple-Alford and Marmurek (1999) claimed that the semantic priming effects generated by the model when simulating the word naming task were due to interference rather than facilitation. To support this claim, they used and unprimed presentation of a target word as a baseline condition. Dalrymple-Alford and Marmurek found that a related semantic prime did not produce shorter simulated naming latencies in the model than did the target-alone baseline condition and sometimes the related prime condition even produced interference relative to this baseline. At the same time, however, unrelated primes generated longer simulated naming latencies than related primes. It was concluded that semantic primes, whether related or unrelated, produce only interference effects. More generally, Dalrymple-Alford and Marmurek argued that attractor networks with fully interconnected units, such as the distributed memory model, could not generate facilitative priming in which one set of units (i.e., semantic) improves processing in another set (i.e., phonological).

In response to this claim regarding attractor networks, I showed that the lack of facilitative semantic priming when the target-alone condition is used as the baseline depends on what happens to the network between trials (Masson, 1999). In their simulations (as in earlier simulations my colleagues and I have conducted), Dalrymple-Alford and Marmurek reset the network to a random pattern of activation after each trial. When the patterns of activation in the network's semantic and phonological modules are random, units in those modules send little, if any coherent, input to units that are selected for updating. Rather, they send a random pattern of input that usually has no systematic influence on the updating of a unit. As units within a module take on a more coherent pattern of activation, they send a stronger, more coherent input signal to other units in the network, causing the network to move into a stable pattern of activation that corresponds to a learned word. One way of measuring the coherence of the pattern of activation within a module, and hence the coherence of the input it sends to other units, is to compute the energy within that module's units. This can be done using the formula for energy shown in the Appendix, but applying it only to the units and connection weights within the module in question.

When the energy of units in the semantic and in the phonological modules is computed, it becomes apparent why a related prime generally fails to produce shorter simulated naming latencies than the target-alone baseline condition. At the start of a target-alone trial, the network is in a random state and energy within the semantic and phonological modules is close to zero. On a primed trial, however, at the onset of the target the network already has moved into a pattern of activation corresponding to the prime that has just been processed. Therefore, the energy in the semantic and phonological modules has grown to some large negative value, depending on how long the prime was processed. Those modules consequently send coherent input signals to the network's units in support of the prime's pattern of activation. When the prime is unrelated to the target, these coherent input signals are not compatible with the target's pattern and they conflict with the input signals generated by the orthographic units when the target pattern is presented. This conflict slows the network's progress into the target's pattern of activation. As Figure 5a shows, this effect increases as a function of energy in the semantic units at the time the target's orthographic pattern is presented to the network. Large negative energy values indicate a more complete instantiation of the prime's pattern of activation. Data in Figure 5a were obtained by presenting an unrelated prime for varying durations and computing energy in the semantic units at the moment the target was presented and computing the number of updating cycles needed for the phonological units to settle on the target's phonological pattern. A relation similar to that shown in Figure 5a also holds for the energy in the phonological units, but in that case it is a weaker relation because there is a relatively small number of phonological units so their state at target onset has a smaller effect on the network's behavior.

Insert Figure 5 about here

When the prime is semantically related to the target, the input of semantic units to phonological units is somewhat compatible with the target pattern as well, leading to a reduction or elimination of the interference effect. The net result is a priming effect as assessed by the difference between related and unrelated prime conditions. The advantage of semantic relatedness, however, is not great enough to generate a facilitation effect relative to the target-alone condition in which there is no competing source of coherent input to semantic and phonological units.

The distributed memory model is capable of producing a facilitation effect relative to the target-alone condition, however, if the network is not reset to a random state at the start of each trial. By starting each trial with the network in the state it reached at the end of the previous trial, even on target-alone trials the network's semantic and phonological modules are in coherent patterns of activation (energy values are large and negative) at the time the target is presented. Under these conditions, Masson (1999) found that a related prime led to shorter simulated naming latency than the target-alone condition and that this facilitation effect increased as a function of how long the prime was processed, at least up to a point (see Figure 5b). With longer prime durations, the facilitation effect began to decrease somewhat because the network moved very deeply into the prime's basin of attraction. As can be seen in Figure 5b, the unrelated primes also produce a facilitation effect at the shortest prime durations. Although this effect seems counterintuitive, in the absence of data from experiments with human subjects it can be taken as an interesting prediction of the model.

Other Types of Primes

As an attractor network, the distributed memory model is designed specifically to react to external input by settling in the nearest basin of attraction. A disadvantage of this "single-mindedness" is revealed in the model's response to three types of primes that differ from the typically used related and unrelated primes. First, Dalrymple-Alford and Marmurek (1999) showed that a neutral prime, coded as a random pattern of activation in the orthographic units, had a systematic effect on the distributed memory model's behavior. By starting the semantic and phonological units in a random state, then presenting the network with a random orthographic pattern, the network will eventually move into a pattern of activation corresponding to the word whose pattern was most similar to the random starting state of the network. Thus, using a random orthographic pattern as a neutral prime did not cause the network to remain in some indeterminate state. Similarly, Bourassa and Besner (1998) used the model to simulate the influence of nonword primes that were orthographically very similar to learned words. In the related prime condition, the nonword prime was orthographically similar to a learned word that was semantically similar to the target word; in the unrelated prime condition the nonword prime was orthographically similar to a word that was semantically unrelated to the target (actual counterparts might be deg-CAT vs. deg-SKY). They found that the model produced a semantic priming effect when this type of prime was used and that the size of this effect increased with longer prime durations, just as occurred for word primes. In contrast, they found that in experiments with human subjects nonword primes generated a priming effect only when they were presented briefly and masked. They found no priming effect with nonwords when the primes were clearly visible to subjects. Finally, some priming experiments with human subjects have used a single word as a neutral prime (e.g., blank or ready). By presenting the chosen word frequently, always followed by an unrelated target, subjects came to react differently to such a prime than to unrelated primes seen only once (e.g., McNamara, 1994; Ratcliff & McKoon, 1988, 1995).

Some additional mechanism(s) would be required for the network to react in a more noncommittal manner to neutral primes, to be unresponsive to clearly presented nonword primes, and to learn to respond differently to primes that are repeatedly presented with unrelated targets. One approach, proposed by Bourassa and Besner as a means of simulating the nonword priming effect, would be to add a verification process in which the familiarity of an orthographic pattern is assessed. If the pattern is found to be unfamiliar, the network would deconstruct the pattern of activation in semantic units created by that pattern.

An alternative and more general idea is to place a constraint on updating semantic and phonological units. Before an orthographic pattern of activation exerts an influence on semantic and phonological units, an assessment is done regarding that pattern's familiarity. When a familiar pattern is detected (as indicated, for example, by a large negative energy value in the orthographic units), semantic and phonological units begin updating. The thoroughness of this assessment would depend on the availability of orthographic information. With a brief, masked presentation, orthographic information would be minimal and the assessment might be preempted or curtailed, allowing a word-like stimulus to influence semantic and phonological units. A gating mechanism of this sort might also be used to prevent a prime word that is repeatedly presented with unrelated targets from influencing semantic and perhaps phonological units, akin to a semantic satiation effect (e.g., Esposito & Pelton, 1971). A learning mechanism would have to be incorporated as well, however, that would be sensitive to the events in which particular words take part. In general, it appears likely that a more accurate simulation of the effects of various neutral and nonword primes will require the implementation of a mechanism that influences and perhaps gates the reaction of semantic and phonological units to orthographic input and perhaps even to input from themselves.

Stimulus Coding

The final issue to be considered is of general concern in models that are intended to explain responses to externally presented stimuli. That is, how should one represent stimulus input in a model? In many models of word reading, the typical approach has been to finesse this question by taking as the starting point an orthographic representation. Although this approach has the advantage of allowing models to move quickly to issues of higher level processing (such as phonological recoding and semantic processing), it has two serious drawbacks. First, it constrains in arbitrary ways the possible means of simulating the effects on stimulus processing that would occur prior to the formation of an orthographic code, such as stimulus degradation or peripheral neurological disorders. Second, and more important, there may be a high price associated with the convenience of adopting such coding schemes. These coding schemes typically are crafted with downstream processing in mind, as in the conversion of orthographic representations to phonological representations that motivated the Plaut et al. (1996) modifications to the Seidenberg and McClelland (1989) model. Empirical evidence is beginning to emerge that challenges the typical implementation of a word's visual form as a linear orthographic string (Tainturier & Caramazza, 1996).

Rather than using consequences for higher order coding as the primary factor in determining how stimuli ought to be coded, significant attention should be directed toward honoring the constraints of perceptual processing and toward how an adaptive system discovers which features of its environment are important in rendering an internal description of that environment (see Rumelhart & Zipser, 1985, for an approach to this problem). By working with an internal coding scheme that ignores this set of constraints, we risk moving model development along seriously misguided trajectories. The argument here is that very careful consideration needs to be given to alternative methods of coding stimulus events and to the differences in subsequent processing implied by these alternative methods.

Conclusion

The distributed memory model described in this chapter has been a useful tool in efforts to account for the interaction between orthographic, phonological, and semantic knowledge. It's simple representational and learning assumptions have permitted an early glimpse of the effects that semantic knowledge can have on word reading tasks that are the foundation of our empirical knowledge concerning word reading processes. At the same time, however, the model's shortcomings with respect to addressing the subtleties of word characteristics that have dominated much theorizing in other quarters (e.g., Besner, this volume; Besner, Twilley, McCann, & Seergobin, 1990; Plaut, this volume; Plaut & Shallice, 1993; Plaut et al., 1996) has made it difficult for the distributed memory model to make contact with that body of literature. In future development of the model, emphasis will be placed on the construction of more principled representations of knowledge and on methods of governing the network's task performance to achieve the flexibility of processing that is so clearly evident in word reading performance.

References

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

Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. Journal of Experimental Psychology: Human Perception and Performance, 10, 340-357.

Becker, C. A., & Killion, T. H. (1977). Interaction of visual and cognitive effects in word recognition. Journal of Experimental Psychology: Human Perception and Performance, 3, 389-401.

Besner, D., & Johnston, J. C. (1989). Reading and the mental lexicon: On the uptake of visual information. In W. Marslen-Wilson (Ed.), Lexical representation and process (pp. 291-316). Cambridge, MA: MIT Press.

Besner, D., & Smith, M. C. (1992). Models of visual word recognition: When obscuring the stimulus yields a clearer view. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 468-482.

Besner, D., Twilley, L., McCann, R. S., & Seergobin, K. (1990). On the association between connectionism and data: Are a few words necessary? Psychological Review, 97, 432-446.

Bodner, G. E., & Masson, M. E. J. (1997). Masked repetition priming of words and nonwords: Evidence for a nonlexical basis for priming. Journal of Memory and Language, 37, 268-293.

Borowsky, R., & Besner, D. (1991). Visual word recognition across orthographies: On the interaction between context and degradation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 272-276.

Borowsky, R., & Masson, M. E. J. (1996). Semantic ambiguity effects in word identification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 63-85.

Bourassa, D. C., & Besner, D. (1998). When do nonwords activate semantics? Implications for models of visual word recognition. Memory & Cognition, 26, 61-74.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing.

Psychological Review, 82, 407-428.

Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. Psychological Review, 100, 589-608.

Dalrymple-Alford, E. C., & Marmurek, H. H. C. (1999). Semantic priming in fully recurrent network models of lexical knowledge. Journal of Experimental Psychology: Learning, Memory, and Cognition, 25, 758-775.

Daneman, M., Reingold, E. M., & Davidson, M. (1995). Time course of phonological activation during reading: Evidence from eye fixations. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 884-898.

Dannenbring, G. L., & Briand, K. (1982). Semantic priming and the word repetition effect in a lexical decision task. Canadian Journal of Psychology, 36, 435-444.

Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. Journal of Memory and Language, 27, 429-446.

Esposito, N. J., & Pelton, L. H. (1971). Review of the measurement of semantic satiation. Psychological Bulletin, 75, 330-346.

Farah, M. J. (1989). Semantic and perceptual priming: How similar are the underlying mechanisms? Journal of Experimental Psychology: Human Perception and Performance, 15, 188-194.

Fodor, J. (1983). The modularity of mind. Cambridge, MA: MIT Press.

Forster, K. I. (1976). Accessing the mental lexicon. In R. J. Wales & E. W. Walker (Eds.), New approaches to language mechanisms (pp. 257-287). Amsterdam: North-Holland.

Forster, K. I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 680-698.

Hebb, D. O. (1949). The organization of behavior. New York: Wiley.

Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. Psychological Review, 98, 74-95.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Science, 79, 2554-2558.

Jacobs, A. M., & Grainger, J. (1992). Testing a semistochastic variant of the interactive activation model in different word recognition experiments. Journal of Experimental Psychology: Human Perception and Performance, 18, 1174-1188.

Jared, D., & Seidenberg, M. S. (1991). Does word identification proceed from spelling to sound to meaning? Journal of Experimental Psychology: General, 120, 358-394.

Johnson, N. F., & Pugh, K. R. (1994). A cohort model of visual word recognition. Cognitive Psychology, 26, 240-346.

Joordens, S., & Besner, D. (1992). Priming effects that span an intervening unrelated word: Implications for models of memory representation and retrieval. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 483-491.

Joordens, S., & Besner, D. (1994). When banking on meaning is not (yet) money in the bank: Explorations in connectionist modeling. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 1051-1062.

Joordens, S., Masson, M. E. J., & Besner, D. (1995, November). Connectionist models and additive effects: Are distinct stages of processing necessary? Paper presented at the annual meeting of the Psychonomic Society, Los Angeles, CA.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. Psychological Review, 87, 329-354.

Lesgold, A. M., & Perfetti, C. A. (Eds.). (1981). Interactive processes in reading. Hillsdale, NJ: Erlbaum.

Lukatela, G., & Turvey, M. T. (1994). Visual lexical access is initially phonological: 1. Evidence from associative priming by words, homophones, and pseudohomophones. Journal of Experimental Psychology: General, 123, 107-128.

Massaro, D. W. (1979). Letter information and orthographic context in word perception. Journal of Experimental Psychology: Human Perception and Performance, 5, 595-609.

Massaro, D. W. (1989). Testing between the TRACE model and the fuzzy logical model of speech perception. Cognitive Psychology, 21, 398-421.

Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), Basic processes in reading: Visual word recognition (pp. 233-263). Hillsdale, NJ: Erlbaum.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 3-23.

Masson, M. E. J. (1999). Semantic priming in a recurrent network: Comment on Dalrymple-Alford and Marmurek (1999). Journal of Experimental Psychology: Learning, Memory, and Cognition, 25, 776-794.

Masson, M. E. J., & Borowsky, R. (1998). More than meets the eye: Context effects in word identification. Memory & Cognition, 26, 1245-1269.

Masson, M. E. J, & Sala, L. S. (1978). Interactive processes in sentence comprehension and recognition. Cognitive Psychology, 10, 244-270.

McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. Psychological Review, 86, 287-330.

McClelland, J. L. (1985). Putting knowledge in its place: A scheme for programming parallel processing structures on the fly. Cognitive Science, 9, 113-146.

McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. Cognitive Psychology, 23, 1-44.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. Psychological Review, 88, 375-407.

McNamara, T. P. (1992). Theories of priming: I. Associative distance and lag. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 1173-1190.

McNamara, T. P. (1994). Theories of priming: II. Types of primes. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 507-520.

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

McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. Journal of Experimental Psychology: General, 126, 99-130.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (November 1972). Activation of lexical memory. Paper presented at the annual meeting of the Psychonomic Society, St. Louis, MO.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1975). Loci of contextual effects on visual word recognition. In P. M. A. Rabbitt & S. Dornic (Eds.), Attention and Performance V (pp. 98-118). San Diego, CA: Academic Press.

Morton, J. (1969). Interaction of information in word recognition. Psychological Review, 76, 165-178.

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.

Neely, J. H. (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.

Norris, D. (1995). Signal detection theory and modularity: On being sensitive to the power of bias models of semantic priming. Journal of Experimental Psychology: Human Perception and Performance, 21, 935-939.

Paap, K. R., & Noel, R. W. (1991). Dual-route models of print and sound: Still a good horse race. Psychological Research, 53, 13-24.

Perfetti, C. A., & Bell, L. (1991). Phonemic activation during the first 40 ms of word identification: Evidence from backward masking and masked priming. Journal of Memory and Language, 30, 473-485.

Perfetti, C. A., Bell, L., & Delaney, S. (1988). Automatic phonetic activation in silent word reading: Evidence from backward masking. Journal of Memory and Language, 27, 59-70.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. Psychological Review, 103, 56-115.

Plaut, D. C., & Shallice T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. Cognitive Neuropsychology, 10, 377-500.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), Information processing and cognition (pp. 55-85). Hillsdale, NJ: Erlbaum.

Postman, L., & Keppel, G. (Eds.). (1970). Norms of word association. New York: Academic Press.

Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. Psychological Review, 95, 385-408.

Ratcliff, R., & McKoon, G. (1994). Retrieving information from memory: Spreading-activation theories versus compound-cue theories. Psychological Review, 101, 177-184.

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

Rayner, K., & Frazier, L. (1989). Selection mechanisms in reading lexically ambiguous words. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 779-790.

Rhodes, G., Parkin, A. J., & Tremewan, T. (1993). Semantic priming and sensitivity in lexical decision. Journal of Experimental Psychology: Human Perception and Performance, 19, 154-165.

Rueckl, J. G., & Oden, G. C. (1986). The integration of contextual and featural information during word identification. Journal of Memory and Language, 25, 445-460.

Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), Attention and performance VI (pp. 573-603). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. Cognitive

Science, 9, 75-112.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. Psychological Review, 96, 523-568.

Seidenberg, M. S., & Tanenhaus, M. K. (1986). Modularity and lexical access. In I. Gopnik & M. Gopnik (Eds.), From models to modules: Studies in cognitive science from the McGill workshops (pp. 135-157). Norwood, NJ: Ablex.

Sharkey, N. E. (1989). The lexical distance model and word priming. In G. M. Olson & E. E. Smith (Eds.), Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (pp. 860-867). Hillsdale, NJ: Erlbaum.

Sharkey, N. E. (1990). A connectionist model of text comprehension. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), Comprehension processes in reading (pp. 487-514). Hillsdale, NJ: Erlbaum.

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

Stolz, J. A. & Neely, J. H. (1995). When target degradation does and does not enhance semantic context effects in word recognition. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 596-611.

Tainturier, M.-J., & Caramazza, A. (1996). The status of double letters in graphemic representations. Journal of Memory and Language, 35, 53-73.

Thompson-Schill, S. L., Kurtz, K. J., & Gabrieli, J. D. E. (1998). Effects of semantic and associative relatedness on automatic priming. Journal of Memory and Language, 38, 440-458.

Van Orden, G. C., Johnston, J. C., & Hale, B. L. (1988). Word identification in reading proceeds from spelling to sound to meaning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 371-385.

Ziegler, J. C., & Jacobs, A. M. (1995). Phonological information provides early sources of constraint in the processing of letter strings. Journal of Memory and Language, 34, 567-593.

Acknowledgement

Preparation of this report was supported by a research grant from the Natural Sciences and Engineering Research Council of Canada. I am grateful to Derek Besner, Ron Borowsky, and Steve Joordens for their many contributions to the collaborative research reported here and to Derek Besner for comments on an earlier version of this chapter.

Appendix

Implementation of the Distributed Memory Model

The learning rule used in the distributed memory model is a version of the Hebb (1949) learning rule, in which connection weights between units are adjusted as a function of the activation states of the two units when a target pattern is instantiated in the network. Each time a pattern is presented for a training trial, the connection weight between a pair of units is changed according to the formula

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

where w_{ij} is the connection weight between units *i* and *j*, and n_i and n_j are the activation states for those units when a target pattern is instantiated. For example, if the pattern of activation for a given word calls for two particular orthographic units to take on the value of 1, the connection weight between them will increase by (1)(1) = 1.

During pattern completion, part of a learned pattern is instantiated in the network while the remaining units in the network are set to randomly determined states. The units that represent the partial pattern are clamped (i.e., fixed at their assigned activation states) and the remaining units are sampled and their activation states updated. A unit's activation state is updated by computing the net input coming into that unit from all other units in the network, according to the formula

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

where net_i represents the net input received by unit *i*. A nonlinear activation function is used to convert the net input into the unit's activation state of 1 or –1. The function is the threshold function

$$\text{if } net_i > 0, \text{ then } n_i = 1, \text{ else } n_i = -1.$$

This updating rule normally causes the pattern of activation across the network's units to move into a stable pattern of activation matching the target pattern that was partially instantiated.

The movement of the network into a learned pattern of activation is a form of gradient descent into a basin of attraction that can be described quantitatively as moving the network into a local minimum of an energy function, *E*, defined as

$$E = -\sum_{i>j} w_{ij}\, n_i\, n_j\ .$$

In a three-dimensional analogy, *E* corresponds to depth in the landscape and has a local minimum at the bottom of a basin of attraction (corresponding to a learned pattern). Each time a unit is updated and its activation state computed, the network has the opportunity to move further down the basin it has entered, reducing *E*. Whenever the updating of a unit causes its state to change, *E* is reduced. On some occasions, the input coming into a unit and the threshold function dictate that a unit take on the same activation value it already holds (i.e., the unit does not change state). In these cases, changing the unit's state of activation would cause *E* to increase. Thus, when the network eventually enters a pattern of activation that corresponds to learned pattern, the updating rule will call for none of the units to change state (i.e., *E* will have reached a local minimum). We refer to this situation as reaching a stable state.

As an alternative to energy, one can use an inverse of that measure generally known as harmony. The simulations of masked priming described here used a particular version of this measure, <u>scaled</u> <u>harmony</u> (Joordens et al., 1995). This measure is a scaled inverse of energy,

$$H_S = \sum_{i>j} c_{ij}\, w_{ij}\, n_i\, n_j\ ,$$

where c_{ij} is a constant equal to the reciprocal of the number of connection weights of the same type as w_{ij} (e.g., semantic unit to semantic unit, or orthographic unit to phonological unit). The scaled harmony measure adjusts the contributions of the three different modules of the network and the connections between them so that all three modules have an equal influence on harmony, independently of differences in number of units in the modules. Without the scaling constant, modules with more units play a greater role in determining the value of harmony. By scaling the w_{ij} values to take into account the number of connection weights of each type (both within and across modules), the potentially arbitrary differences in number of units in the modules will not affect the harmony measure. Although the decision to allow each module to contribute equally to harmony is also arbitrary, it was deemed a reasonable starting assumption. Variations in task demands and materials may affect the validity of this assumption. The maximum value that scaled harmony can take on is equal to the total number of learning trials, and is attained only if exactly the same pattern is learned on each trial. As different patterns are learned, the Hebbian learning rule leads to connection weights that vary in strength, reducing the maximum value of scaled harmony the network can attain.

Table 1

Mean Response Latency and Error Percentage as a Function of Lexical Ambiguity

	Response Latency		% Error	
Data source and task	Ambig.	Unambig.	Ambig.	Unambig.
Results from human subjects				
Word naming	494	495	4.1	5.1
Lexical decision				
Pronounceable nonwords	637	647	3.2	4.3
Consonant string nonwords	569	567	2.3	2.6
Results from simulations				
Word naming	291	290	3.7	3.7
Lexical decision				
Pronounceable nonwords	116	119	3.2	5.2
Consonant string nonwords	54	54	3.6	4.0
Gaze duration during comprehension	689	629	0.8	0.0

Note. Adapted from Borowsky and Masson (1996). Latencies for human data are in milliseconds and latencies for simulated data are in cycles.

Figure Captions

Figure 1. Architecture of the distributed memory model. Each type of word knowledge is represented as a pattern of activation across a set of processing units. Arrows indicate flow of activation within and between sets of units.

Figure 2. Activation over time of orthographic, phonological, and semantic knowledge about a word whose orthographic pattern is presented to the network (a). Activation is measured as the proportion of units in a module that are in a state that corresponds to target word's pattern. Activation of target word knowledge within each module is compared with activation of knowledge about an orthographically related, a semantically related, and an unrelated word that occurs as a byproduct of processing the target word (b, orthographic module; c, phonological module; d, semantic module).

Figure 3. Mean cycles to settle phonological units under conditions of clear and degraded visual input. The effect of semantic priming is increased when degraded input is used. Error bars indicate the within-subjects confidence interval for the means.

Figure 4. Growth of scaled harmony over processing cycles (a) and mean cycles to reach harmony criterion (b) for the simulation of masked identity priming in lexical decision. The harmony criterion for making a positive lexical decision was set at .15, as shown in section a. The error bars in section b represent the within-subjects confidence interval for the means.

Figure 5. Relation between energy in semantic units at target onset and cycles to settle on the target's phonological pattern when an unrelated prime is presented (a) and mean cycles to settle on a target's phonological pattern as a function of prime condition when the network is not reset to a random state at the start of each trial (b).

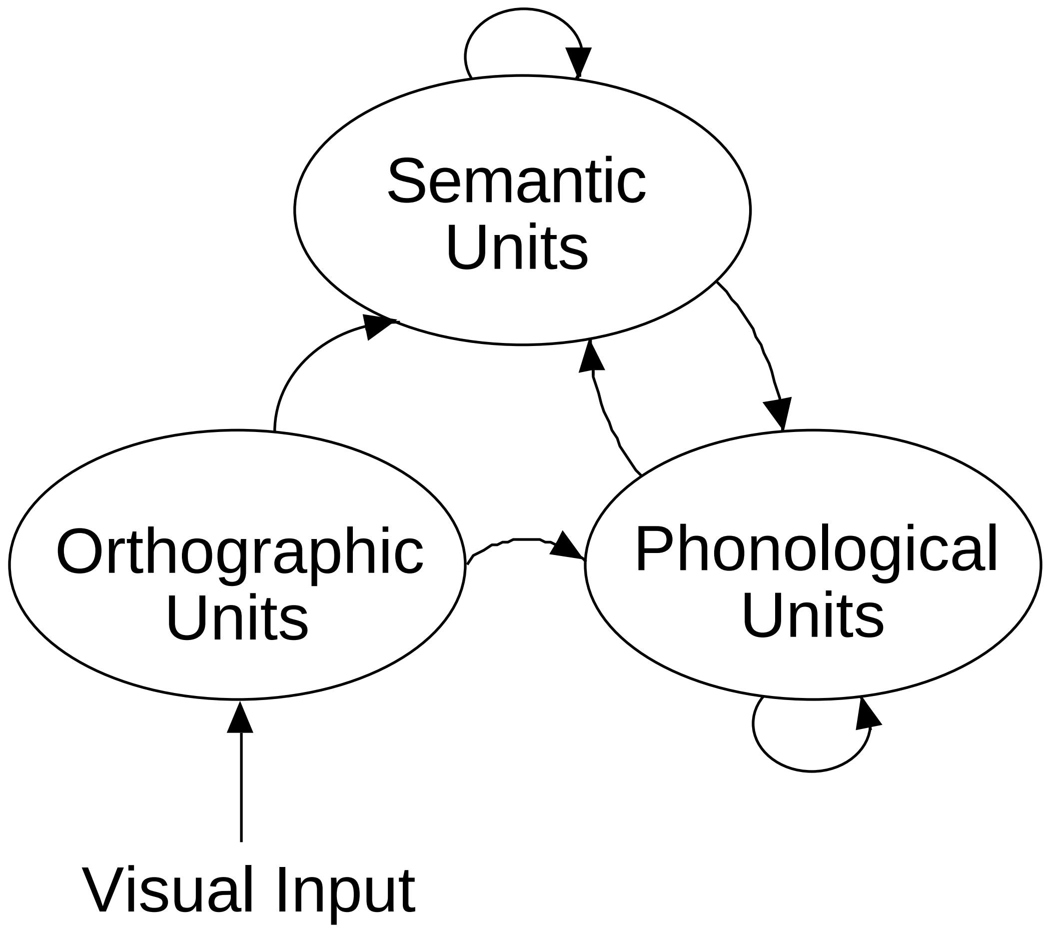
Semantic
Units
Orthographic
Units
Phonological
Units
Visual Input

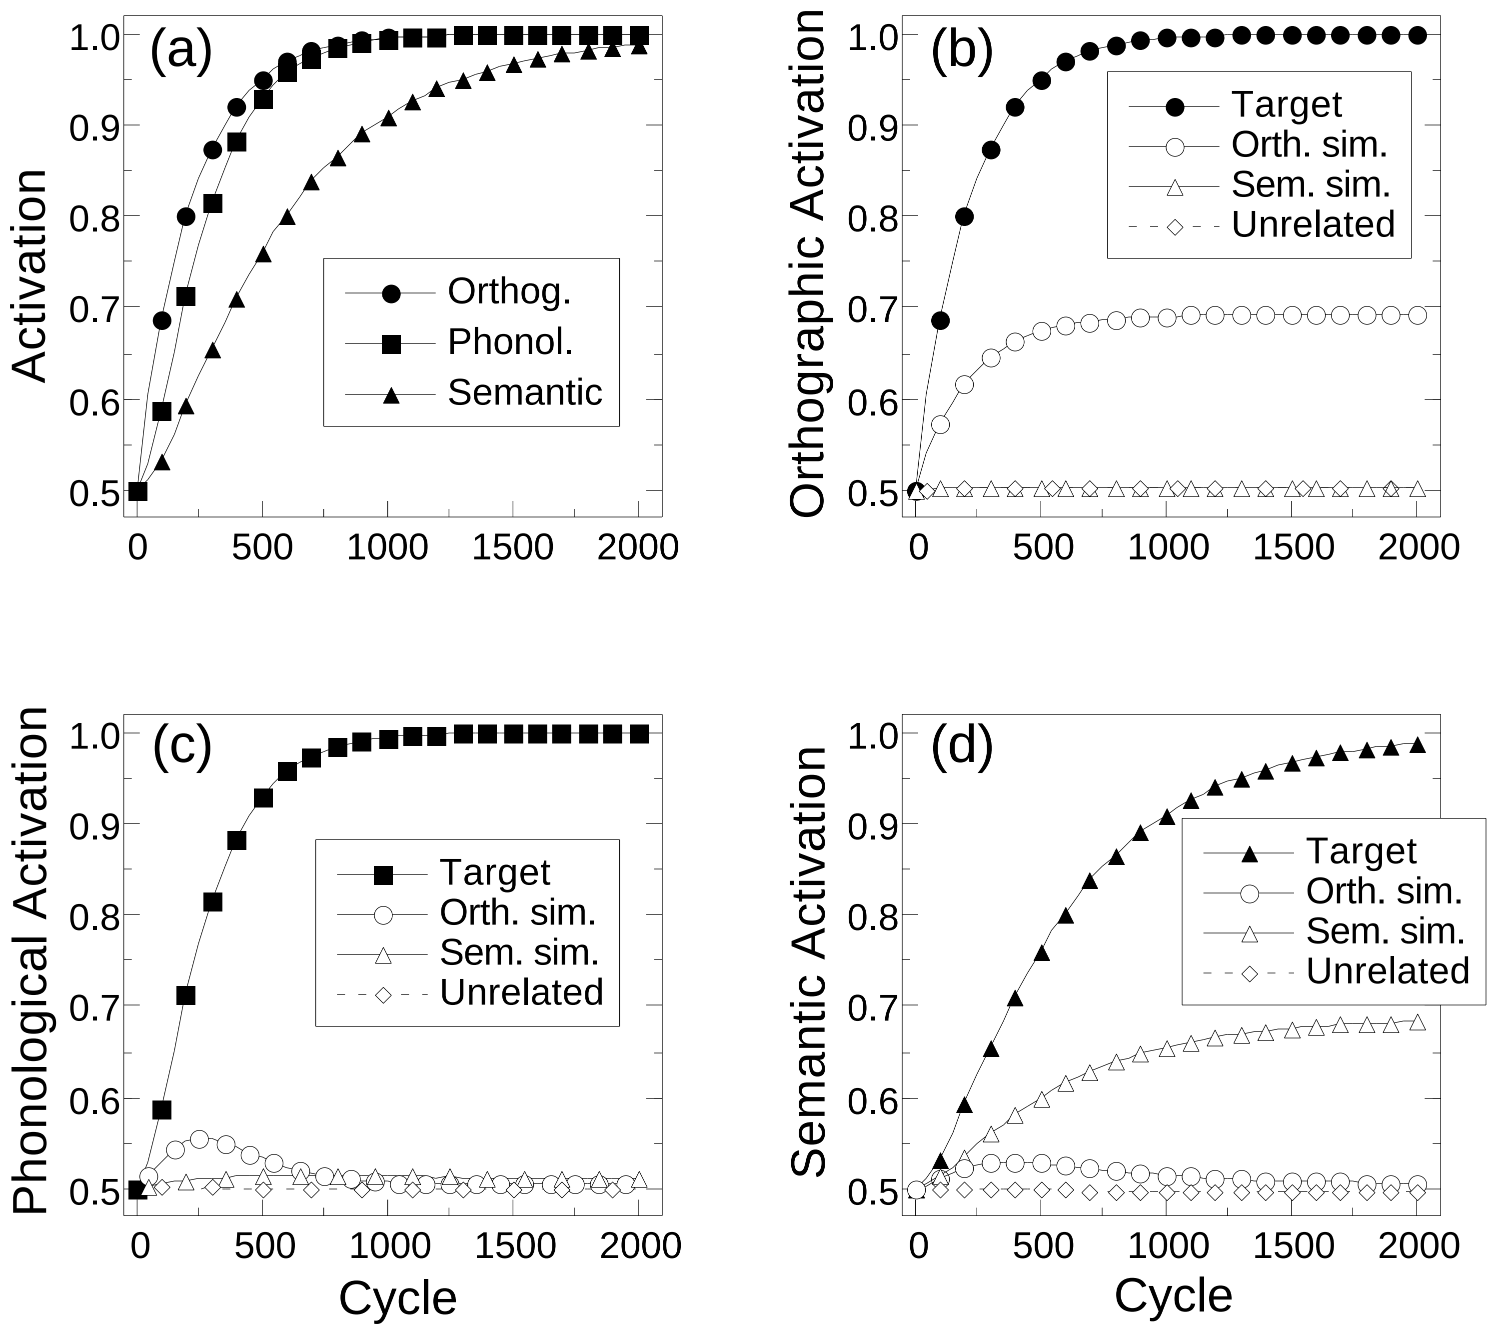

(a)
Activation
Orthog.
Phonol.
Semantic
(b)
Orthographic Activation
Target
Orth. sim.
Sem. sim.
Unrelated
(c)
Phonological Activation
Target
Orth. sim.
Sem. sim.
Unrelated
Cycle
(d)
Semantic Activation
Target
Orth. sim.
Sem. sim.
Unrelated
Cycle

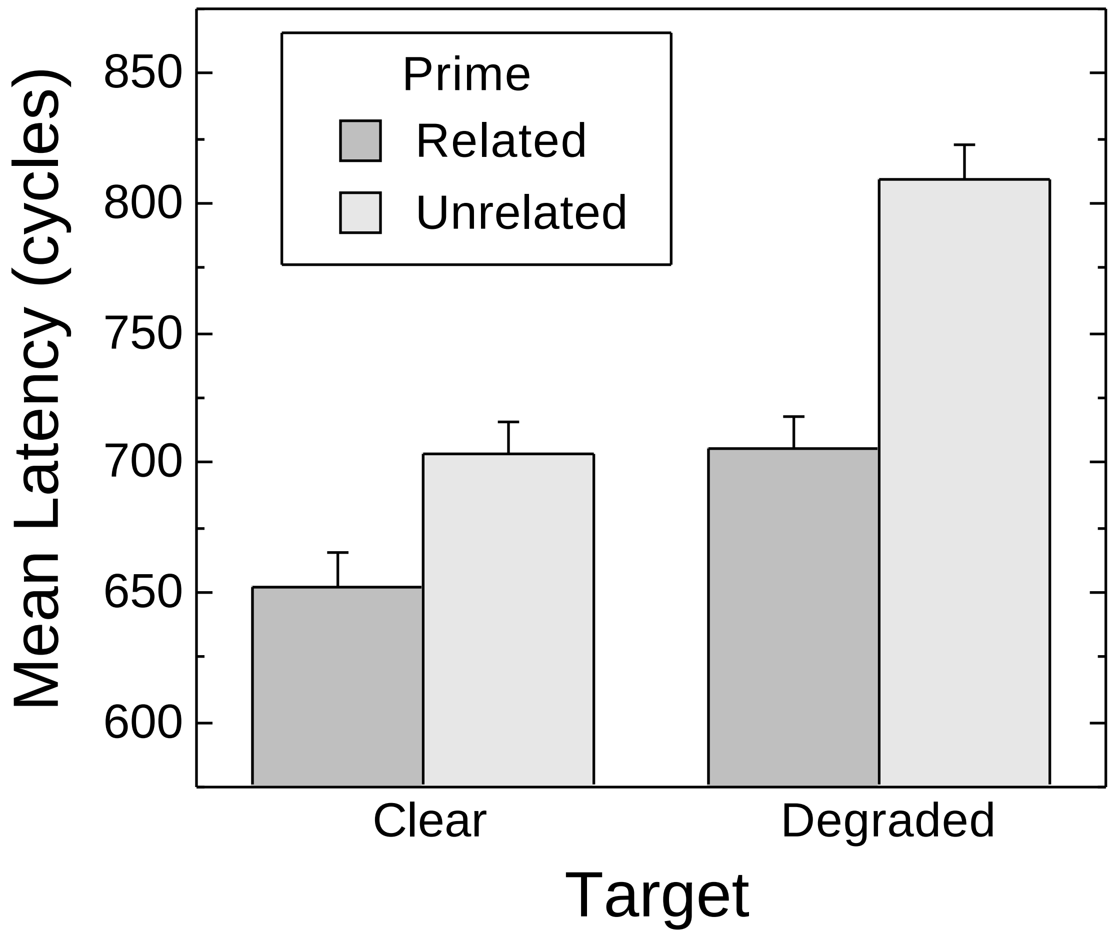
Prime
Related
Unrelated
850
800
750
700
650
600
Mean Latency (cycles)
Clear
Degraded
Target

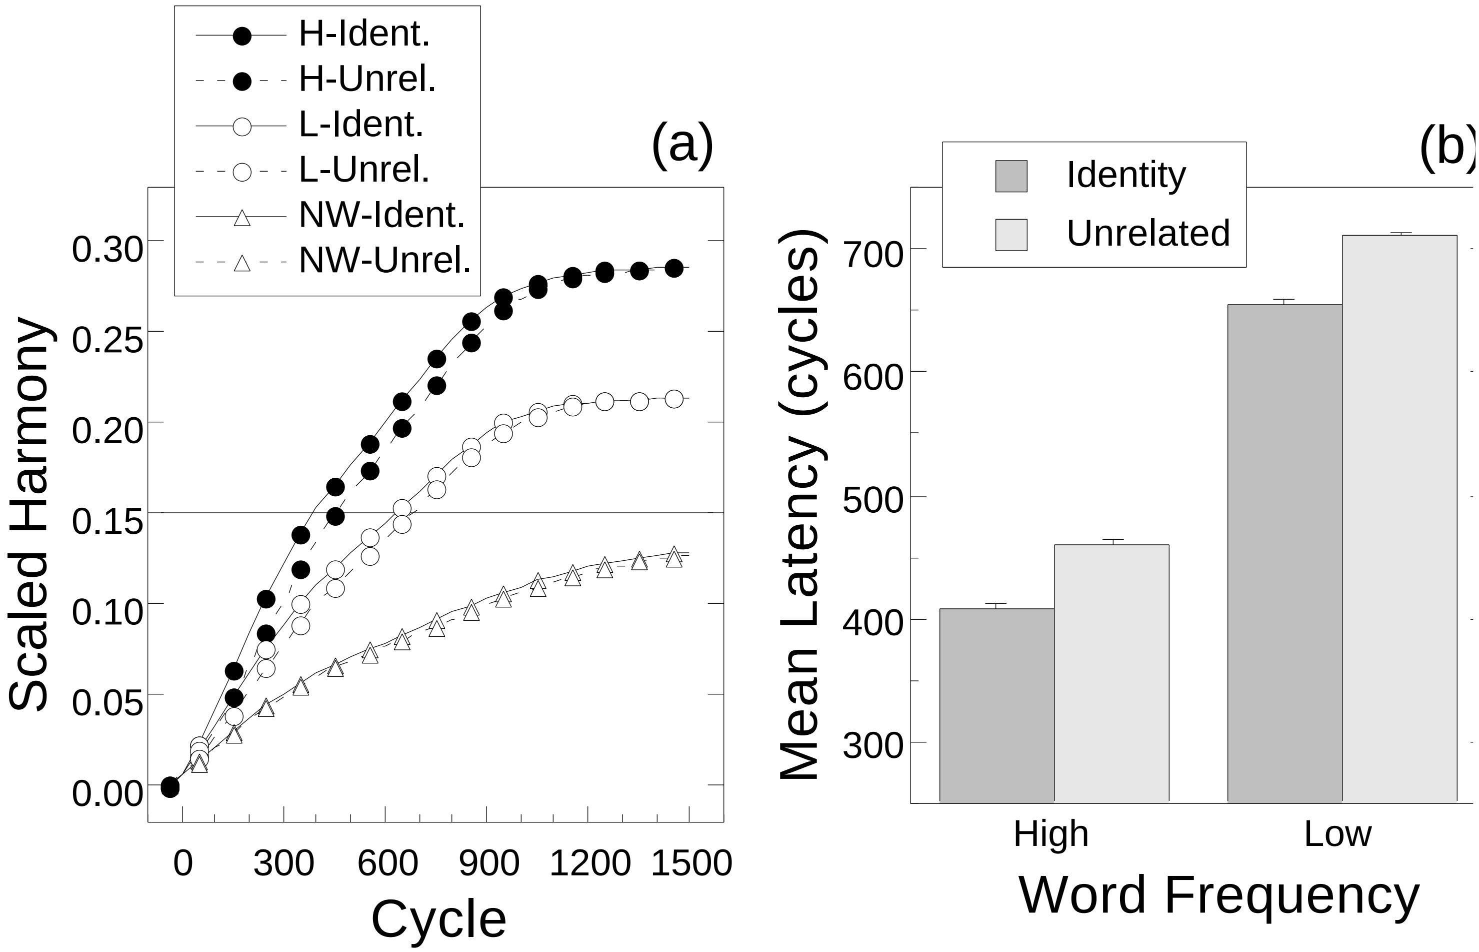
H-Ident.
H-Unrel.
L-Ident.
L-Unrel.
NW-Ident.
NW-Unrel.
(a)
Scaled Harmony
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
300
600
900
1200
1500
Cycle
(b)
Identity
Unrelated
Mean Latency (cycles)
700
600
500
400
300
High
Low
Word Frequency

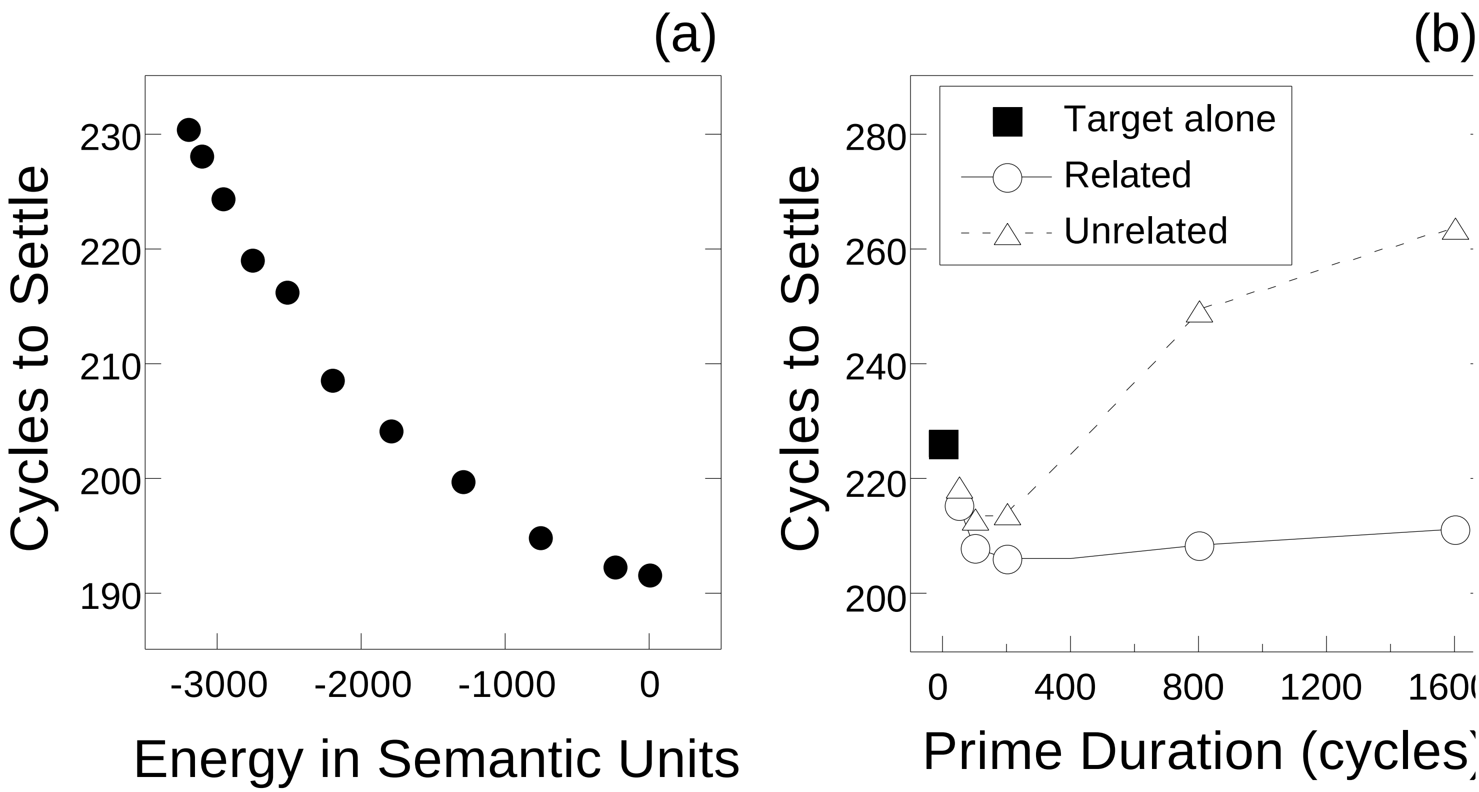

(a)
Cycles to Settle
230
220
210
200
190
-3000
-2000
-1000
0
Energy in Semantic Units
(b)
Target alone
Related
Unrelated
Cycles to Settle
280
260
240
220
200
0
400
800
1200
Prime Duration (cycles