A Connectionist Model of Form-related Priming Effects

Robert R. Peterson, Gary S. Dell, Padraig G. O'Seaghdha
University of Rochester

Abstract

In contrast to the results of many previous studies, Colombo (1986) has demonstrated that form-related priming is sometimes inhibitory. Colombo proposed that inhibition reflects the suppression of lexical items orthographically related to the prime. We suggest, however, that form-related inhibition arises as a result of competition between discrepant prime-target phonemes. During the phonological encoding of the target word, active phonemes from the prime might be mistakenly selected, causing a delay in responding. We present a connectionist model that implements this account, and simulates the empirical data. The model is supported by the results of an experiment that distinguishes between the lexical suppression and phonological competition views.

INTRODUCTION

It has often been demonstrated that the processing of a word can be facilitated by the prior presentation of a formally related word. For example, Meyer, Schvaneveldt, and Ruddy (1974) found shorter lexical decision latencies to words following phonologically and orthographically similar primes (e.g., BRIBE-TRIBE) than to targets following unrelated primes (e.g., FENCE-TRIBE). Recently, however, Colombo (1986) has demonstrated that form-related primes can inhibit responses to high frequency targets. The present paper focusses on the inhibition found by Colombo and her theoretical explanation for the effect. We will suggest an alternative explanation for the inhibition, and will present a connectionist model that accounts for the Colombo data.

The Colombo View: Word-level Inhibition

Colombo suggested that when a prime word is presented, it raises the activation level of a set of letter detectors, and subsequently activates a set of word nodes that are at least partially consistent with those letters. This process results in heightened activation for words orthographically consistent with the prime, and therefore facilitates the recognition of those words. To explain the inhibition found for high frequency targets, Colombo assumes that orthographically similar lexical items inhibit one another, but that this inhibition occurs only for nodes that are highly activated. That is, Colombo argues that lexical nodes have an *inhibition threshold*, and become susceptible to inhibitory influences only when their total activation surpasses this threshold level. Since high frequency words have relatively high resting levels of activation, they quickly surpass their inhibition thresholds. Low frequency words, however, start at such a low level of activation that they never reach their inhibition thresholds. Overall, then, low frequency words receive primarily letter-to-word facilitation, while high frequency words receive initial facilitation followed by word-to-word suppression.

An Alternative View: Phonological Competition

We agree with Colombo's claim that facilitation can arise as a result of activation spreading to lexical nodes that share letters with the prime. We disagree, however, with her explanation of form-related inhibition. While Colombo assumes that inhibition occurs at the lexical level, we argue that it arises at the phonological level instead, resulting from competition between discrepant phonemes of the prime and target. According to our view, it is difficult to respond to the word MAN following the prime FAN, not because the lexical item MAN is inhibited, but rather because the phoneme M in MAN must compete with the already activated F of FAN. Since the presentation of the word MAN will tend to activate the lexical node FAN (due to the shared letters), the F phoneme will be initially quite active, creating the possibility that the F rather than the M phoneme will be selected during the phonological encoding of MAN.

If inhibition is indeed caused by competition at a phonological level, why should the effect appear only with high frequency targets? This effect is likely due to the fact that high frequency words are recognized very quickly. This rapidity in recognition has two consequences, both of which accentuate inhibition. The first consequence is that a high frequency target is likely to get little facilitation from orthographic overlap with the prime, since the baseline recognition rate for the target is already so quick. Therefore, if an inhibitory process exists, it will not be washed out by concurrent facilitative processes.

A fast recognition rate also means that phonological encoding will occur more quickly for high than for low frequency words. Thus, for high frequency targets, selection of the target's component phonemes will occur when the prime's phonemes might still be active, thereby increasing the probability that a phoneme from the prime will be mistakenly selected. Recovering from such an error is likely to be time-consuming, and therefore will lead to an overall inhibition effect. In contrast, the selection problem will be less troublesome for low frequency words, since these words are recognized fairly slowly, allowing activation of the prime's phonemes to decay before selection of the target's phonemes occurs.

We make the following specific assumptions regarding phoneme selection. Retrieval of a lexical item makes available information pertaining to the general phonological form of the word. This information might be represented as an abstract frame which specifies the number, type, and order of the phonemes in the word, as well as the word's syllabic structure and stress pattern (Brown & McNeill, 1966; Dell, 1988; Stemberger, in press). The frame can be thought of as containing slots for each of its phonemes, with the slots being filled by selecting from among activated phonological nodes. For example, retrieval of the word MAN might activate a CVC frame, which specifies that three phonemes are to be retrieved: an initial consonant, a medial vowel, and a final consonant. In our model, high frequency words make their frames available more quickly than do low frequency words, and thus attempt to link phonemes to the frames at an earlier point in processing.

Thus, on our view, there are two qualitatively different effects that a prime word has on an orthographically and phonologically similar target word. First, there is facilitation at the lexical level, due to the activation of words sharing letters with the prime. Second, there is confusion at the phonological level over the incompatible phonemes of the prime and target. This confusion is only problematic, however, if an attempt is made to select the target's phonemes while the prime's phonemes are highly activated. Form-related priming can be seen, therefore, as facilitating lexical retrieval, but potentially interfering with the specification of a word's complete phonological form. We have implemented these ideas within a connectionist framework, and the remainder of this paper is a presentation of our model.

THE PHONOLOGICAL COMPETITION MODEL

Components of the Model

The model is made up of three distinct levels of representation: *letters*, *words*, and *phonemes* (see Figure 1). A node at the letter level corresponds to a given letter in a specific word position. The model has been constructed to process 3-letter words only, so there are three nodes for each letter. Likewise, each node at the phoneme level stands for a given phoneme in a particular word position. All of the words in the model are made up of three phonemes, so a given phoneme is represented three times, once for each word position. For both phoneme and letter nodes, the resting level of activation was set at 0.

A node at the word level corresponds to a single word. In the implementation described here, the model was provided with six words: *CAT*, *CAP*, *CAD*, *PEG*, *PEN*, and *PEZ*. The words *CAT* and *PEG* served as related and unrelated prime words, respectively. *CAP* and *CAD* were the critical target items, with *CAP* serving a high frequency target and *CAD* as a low frequency target. The lexical items *PEN* and *PEZ* were included so that related and unrelated primes would have equivalent lexical neighborhoods. *PEN*, like *CAP*, is a high frequency word, and *PEZ*, like *CAD* is a low frequency word. The resting levels were set at 0 for high frequency words, and -50 for low frequency words. The primes (*CAT* and *PEG*) were given an intermediate resting level of -25.[1]

---

[1] Our assignment of frequency levels to each of the six words does not necessarily correspond to those words' actual frequencies in the language. Our purpose was simply to create high and low frequency lexical nodes that share letters and phonemes. The particular words that we chose to use can best be thought of as convenient labels for these nodes, and hence the actual characteristics of these words in the language is largely irrelevant to our endeavors.

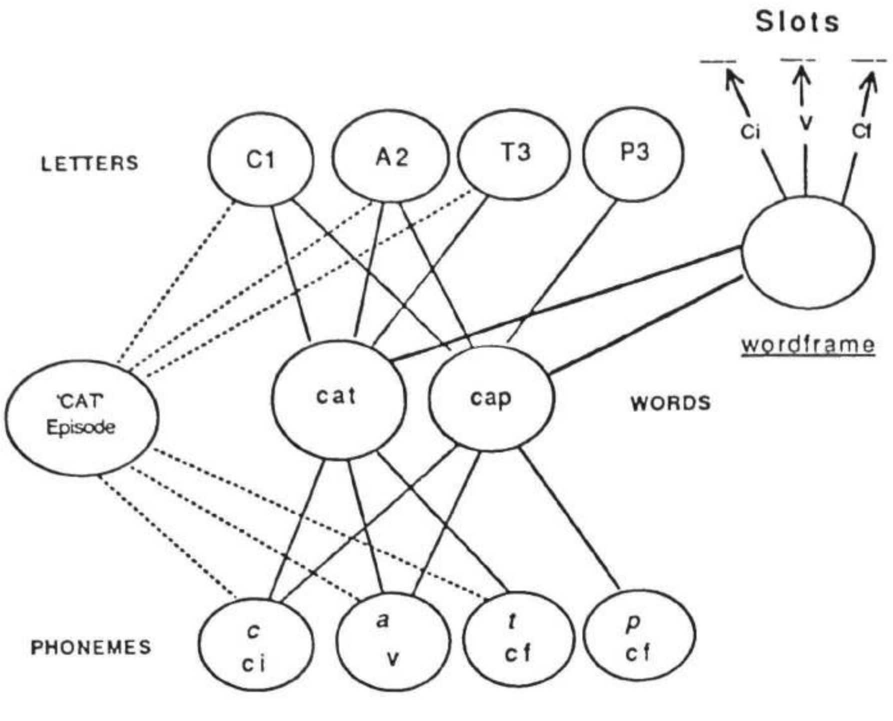

FIGURE 1. Structure of the Model

Connections

There are both excitatory and inhibitory connections in the model. All connections are between nodes at adjacent levels. Word nodes have excitatory connections to their corresponding letter and phoneme nodes (weight = .03). Letter nodes have excitatory connections to words that contain them (weight = .03), and inhibitory connections to other words (weight = .04). Phoneme nodes have excitatory connections to their corresponding words (weight = .03), but no inhibitory connections.

Activation Function

The activation level of a node, at any particular time, is determined by three factors: the node's activation level at the previous timestep, activation received from other nodes during the current timestep, and activation lost during the current timestep due to decay.

The activation received from other nodes is determined according to Equation 1:

$$n_i(t) = \sum_j \alpha_{ij} e_j(t) - \sum_k \gamma_{ik} i_k(t) \tag{1}$$

where $n_i(t)$ is the current input to a node, $e_j(t)$ is the activation of an excitatory neighbor of the node, and $i_k(t)$ is the activation of an inhibitory neighbor of the node. α_{ij} and γ_{ik} are weight constants for excitatory and inhibitory links, respectively.

The amount of decay during a timestep is given in Equation 2:

$$d_i(t) = \Theta_i(a_i(t) - r_i) \tag{2}$$

where $d_i(t)$ is the amount of decay for the node, Θ_i is a constant decay rate (.09 in the present model), $a_i(t)$ is the node's current level of activation, and r_i is the node's resting level of activation. Thus, the amount of decay is proportional to a node's activation relative to its resting level.

The activation of a node at time $t + \Delta t$ is equal, then, to the activation of the node at time t plus the input from other nodes at time t, minus the node's decay. This is expressed mathematically in

Equation 3.

$$a_i(t + \Delta t) = a_i(t) + n_i(t) - d_i(t) \tag{3}$$

There are two important qualifications to Equation 3. First, activation can never go below a node's resting level. Second, we set a maximum level of activation (300) for all the nodes. The output of a node is equal to its activation level if that level is positive. If the node's potential is less than 0, however, the node sends no output.

Episodic Node

In the model, there is an episodic node which operates slightly differently than the rest of the nodes discussed above. The episodic node resides at the lexical level, has a resting level of 0, and has a decay rate of 1.0. During the course of a priming trial, the episodic node establishes connections with the letter and phoneme nodes of the prime, and thereby establishes an episodic memory of the processing of that prime. The actual functioning of the episodic node is described more fully in the next section.

How the Model Works

The model is intended to simulate a priming paradigm, in which a prime word is presented for a certain amount of time and is immediately followed by the presentation of a target word. Prime presentation is simulated by setting the activation level of each of the prime's letter nodes to 300 (their maximum activation). The model is then run for 20 timesteps. The activation of the prime's letter nodes remains fixed at a value of 300 for the entire 20 steps. During these steps, activation spreads throughout the model (from letters to words, from words to both phonemes and letters, and from phonemes to words). The word node corresponding to the prime gets highly activated, primarily due to inputs from its three letter nodes. Orthographically similar words also become active, although to a much lesser extent. These nodes receive excitatory input from the two consistent letter nodes, but receive inhibitory input from the inconsistent third letter.

Following the presentation of the prime, links are created from the prime's letter nodes to the episodic node (weight = .10), and from the episodic node to the prime's phoneme nodes (weight = .015). The episodic node, therefore, is a generic node that is recruited by the model to bind together patterns of activation at the letter and phonological levels. This bound configuration constitutes the model's episodic memory of the prime.

Following the establishment of the episodic links, the presentation of the target is simulated by setting the target's letter nodes to the maximum value of 300. The activation levels of all other letter nodes are set to 0. The model is then run until the target's lexical node reaches its maximum level of activation. Activation spreads through the model in the manner described above for the processing of the prime, with the exception that there are now episodic links active in the model. If the target shares letters with the prime, it will tend to activate the episodic node, which in turn will send activation to the prime's phonemes. Thus, the presentation of an orthographically related target word in effect *reminds* the model of its recent experience with the prime, and causes it to re-create the corresponding phonological representation.

Making Responses

In an actual priming experiment, a subject would be required to make some response to the target. We estimated the model's response time based on Equation 4:

$$RT = \beta l_i + \psi(1 - p(R_i, t)) + \kappa \tag{4}$$

The term βl_i is a measure of lexical access time, with l_i being the number of timesteps required for the target's lexical node to reach its maximum potential and β being a constant specifying the duration of one timestep (5 msec). The component $\psi(1 - p(R_i, t))$ is an estimate of the processing time that is incurred when there is an incorrect selection of the target's critical phoneme. The critical phoneme is the one that is unique to the target, in related prime-target pairs. For example, in the prime-target pair, *CAT-CAD*, the *D* in *CAD* is critical. The probability of correctly selecting the critical phoneme ($p(R_i, t)$ in Equation 4) depends on its activation relative to the activation levels of other phonemes at the same word position. This probability was determined in the same way as in McClelland and Rumelhart (1981), except that it was based on the activation at a single timestep (i.e., the step at which the target's lexical node reached threshold), rather than being based on a running average of a node's activation over time. The probability, then, that an incorrect phoneme is selected is given by $(1 - p(R_1, t))$. The amount of time that is associated with an incorrect selection is given by the constant ψ (set to 150 msec). Finally, κ is a constant (set to 450 msec) which reflects the time required for all of the processing not explicitly represented in the model (e.g., encoding of the stimulus letters, establishment of a motor code for the response, response execution, etc.).

By including the second component in Equation 4, we are assuming that the response to a target word is sensitive to the ease with which the target's phonological form is derived. It is important, therefore, to specify how and when the selection of the target's phonemes occurs. As noted in the introduction, we are assuming that the activation and retrieval of a lexical item makes available an abstract phonological frame that guides the selection process. Although the present model does not contain phonological frames per se, we nevertheless capture the functional effect of these frames by making the selection of phonemes dependent upon the target's lexical node reaching its maximum level of activation.

In the model, there are two important factors that influence the strength of the target's phonemes relative to competing phonemes. The first is how soon phonemes are selected. The earlier the selection process is initiated, the greater the chance that an *incorrect* phoneme is selected. This effect occurs because the prime's phonemes are highly active when the target is first presented, and it takes some time for their activation to decay. A second factor influencing phoneme selection is the relationship between the prime and target. Orthographically and phonologically related prime-target pairs result in greater selection errors than unrelated pairs. This effect occurs because the processing of the related target tends to activate not only the target's lexical node, but also the episodic node and the prime's lexical node. These latter two nodes send activation to the prime's phonemes, thus decreasing the likelihood that the target's critical phoneme will be selected.

Simulations

Relatedness x Frequency Interaction

Colombo (1986) found slower lexical decision latencies to high frequency targets preceded by orthographically and phonologically related primes, relative to unrelated controls. For low frequency targets, on the other hand, the priming effect was facilitatory. We successfully simulated the Colombo results with our model. However, the inhibition effect in our model is not due to direct inhibition of lexical nodes, as proposed by Colombo. Rather the inhibition arises as a result of incorrectly selecting the phonemes of the prime during the processing of the target. This inhibition effect interacts with target frequency in the following way. With high frequency words, the selection process occurs soon after target presentation, when the prime's phonemes are still active. Low frequency words select later, when the activation of the prime's phonemes has subsided. Thus, there is less competition during the selection process for the low frequency targets.

In the simulation, the model was presented with either a high or a low frequency target word (*CAP* or *CAD*, respectively) preceded by either a related or unrelated prime (*CAT* or *PEG*). Figure 2 shows the results of the simulation. The probability of correctly selecting the target's critical phoneme is plotted as a function of the number of timesteps following the target's presentation. In this figure, the four Frequency x Relatedness conditions are plotted separately. The endpoint of each line reflects the timestep at which the target's lexical node reached its maximum activation and phoneme selection occurred.

Several aspects of this figure require comment. First, across timesteps, unrelated targets have a higher probability of correctly selecting their critical phoneme than do related targets. This effect occurs because, on related trials, the target tends to reactivate the prime's phonemes, thus increasing the likelihood of selecting a phoneme from the prime rather than the target. On an unrelated trial, there

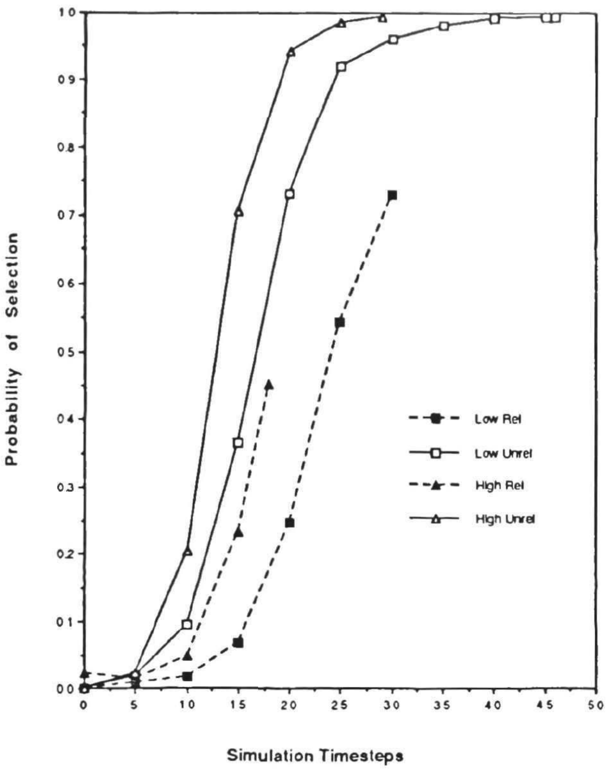

FIGURE 2. Selection of the Critical Phoneme

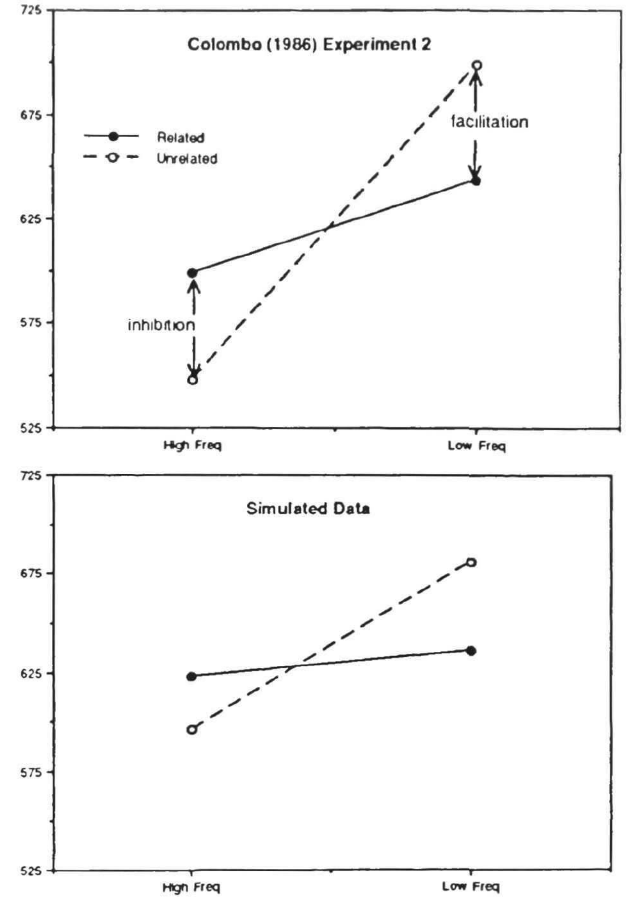

FIGURE 3. Simulation of Relatedness x Frequency Effect

is no overlap in letters between the prime and target, hence the prime's lexical node is inhibited by each of the target's letter nodes. With no lexical support, the activation of the prime's phonemes quickly decays, thereby increasing the probability of correctly selecting the target's critical phoneme.

A second effect shown in the figure is that high frequency targets, at a given timestep, have a higher probability of correct selection than do low frequency targets (this can be seen most clearly by comparing high and low frequency targets within a given level of relatedness). This frequency effect occurs because high frequency words have higher resting levels of activation than do low frequency words and hence activate their component phonemes more quickly.

Of primary interest in this figure, however, is the probability of correct phoneme selection *at the timestep when selection actually occurs* (i.e., the endpoint of each line). At these points, high frequency words have an overall *lower* probability of correct selection than do low frequency words, and this frequency effect interacts with relatedness. For unrelated targets, there is no difference in selection probability for low and high frequency targets (both have reached a ceiling probability of over 99% when selection occurs). For related targets, on the other hand, there is a large advantage for low frequency words (76% probability for the low frequency target, and 45% for the high frequency target). This difference is due to the fact that the high frequency word node reaches its maximum activation more quickly, thereby not giving its critical phoneme enough time to become sufficiently activated.

The response times for each of the four conditions were calculated using Equation 4. The results are presented in Figure 3, along with the results from the Colombo (1986) study. As can be seen, the model's fit to the data is quite good. In both the simulation and Colombo's experiment, there is facilitation for low frequency targets and inhibiton for high frequency targets. Thus, the model accounts for the Colombo data without positing inhibition among lexical candidates.

Varying Episodic Strength

Inhibition occurs in our model as a result of difficulty in forming a phonological representation of the target word. This difficulty arises because the target word tends to activate both the prime's lexical node and the episodic node. By incorporating the episodic node, we are proposing that inhibition depends on the formation of an episodic trace of the prime. We would predict, therefore, that the strength of this trace should influence the size of the inhibition effect. For example, if the prime is presented so briefly that no episodic representation is formed, it follows from our model that inhibition should be substantially reduced. Consistent with this prediction, Forster (1987) has demonstrated that, with subliminal prime presentation, form-related targets are facilitated (see also Humphreys, Evett, & Taylor, 1982).

In order to quantify the relationship between the strength of the episodic trace and phonological priming, we ran a further series of simulations varying the strength of the letter-to-episode connections in the model. We varied the weight of these connections from 0 to .15 (recall that we used a weight of .10 in the previous run of the model). By performing this manipulation, we are changing the extent to which the model is influenced by memories of the prime. A small weight on the letter-to-episode link implies that the model is able to ignore, in a sense, its prior experience with the prime. A large weight suggests that there is a strong tendency to re-create the prime during the processing of the related target.

Except for varying the letter-to-episode connection strength, the model was run as before. Figure 4 shows the results of these simulations. With very weak, or nonexistent episodic traces (specifically weights .025 and 0) there was actually a small facilitation effect for the high frequency target, and a large facilitation effect for the low frequency target. With increases in episodic strength, the high frequency target began to show inhibition, while the low frequency target showed decreasing levels of facilitation. With an episodic weight of .150 (the largest weight used) there was virtually no facilitation for the low frequency target, and a large inhibitory effect for the high frequency target.

There are three different experiments whose data are described reasonably well by these simulations (see Figure 4). The first is the Colombo experiment presented in Figure 3, whose data our simulation captures nicely using a letter-to-episodic weight of .10. It is not surprising that our model does well with the Colombo study, since the model was constructed in large part to reproduce her results. However, in addition to Colombo's results, the model also captures the results of two recent experiments by Lupker and Williams (1987). In one experiment, Lupker and Williams attempted to replicate Colombo,

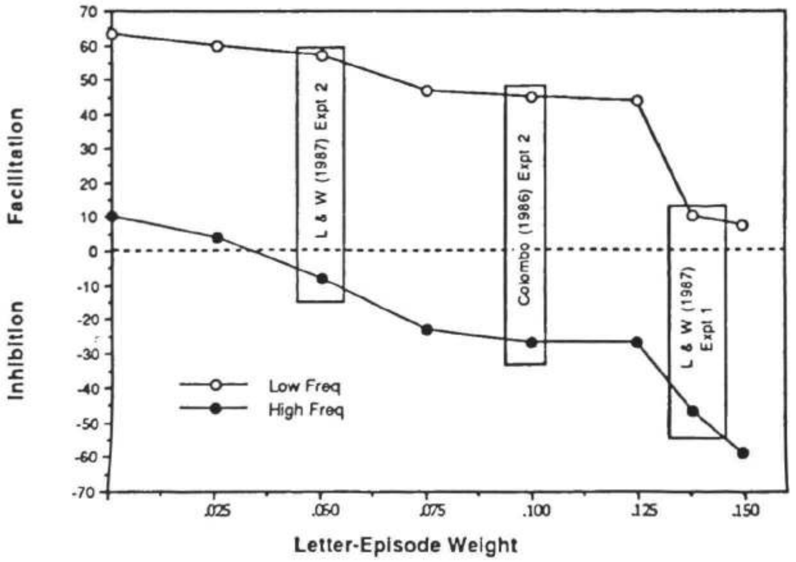

FIGURE 4. Effect of Episodic Strength on Priming

using English rather than Italian materials. They replicated the inhibition effect for high frequency targets, but did not find a significant facilitation effect for low frequency targets. Our model shows a similar pattern of priming with a letter-to-episode weight of .138. In a second experiment, Lupker and Williams used the same materials, but had subjects make lexical decisions to both the prime and the target (in their first experiment, and in Colombo's experiment, subjects made a response to the target only). In this second experiment, there was facilitation for low frequency targets, but only a small (and nonsignificant) inhibition effect for high frequency targets. This pattern of results occurs in our simulations with a weight of about .05.

While the data from the two Lupker and Williams studies and the Colombo experiment may at first appear to be quite disparate, they actually are quite compatible when viewed within the context of Figure 4. According to the analysis provided here, the three experiments differ in terms of vulnerability to episodic interference. Unfortunately, it is difficult to pinpoint the specific aspects of the experiments that might have resulted in these episodic differences. It seems likely, however, that the interference effect might be sensitive to very subtle details within the experimental environment. For example, if the instructions given to the subject at the beginning of the experiment strongly emphasized paying attention to the prime, interference might be more substantial than if the instructions merely informed the subject that a prime word would be presented. Further, the results of Lupker and Williams' second experiment suggest that making an overt response to the prime might significantly modify the nature of the episodic representation of that prime. This modified representation might lead to less confusion during target processing, thus diminishing the interference effect.

CONCLUSIONS

We have presented a model in which there are two distinct loci of form-related priming. Facilitation arises as a result of the activation of orthographically related neighbors during the processing of the prime. On the other hand, inhibition arises during the selection of the target's constituent phonemes, if the prime's phonemes are mistakenly retrieved. Our model is clearly different from that of Colombo, who proposes that both facilitation and inhibition are lexical-level effects. One way to experimentally test our model against Colombo's view is to compare prime-target pairs that are nonhomographic homophones (e.g., *HARE-HAIR*) with pairs that are nonhomophones (e.g., *HATE-HAIR*). Because the primes and targets are orthographically related neighbors in both cases, Colombo would predict a similar pattern of priming for the two types of items. We have conducted this experiment and found inhibition for nonhomophonic items, but facilitation for homophones. This outcome provides strong support for our model: In the absence of phonological competition, orthographically similar words engender facilitation.

REFERENCES

Brown, R., & McNeill, D. (1966). The "Tip of the Tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior, 5*, 325-337.

Colombo, L. (1986). Activation and inhibition with orthographically similar words. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 226-234.

Dell, G. S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. *Journal of Memory and Language, 27*, 124-142.

Forster, K. I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *Quarterly Journal of Experimental Psychology, 39A*, 211-251.

Humphreys, G. W., Evett, L. J., & Taylor, D. E. (1982). Automatic phonological priming in visual word recognition. *Memory & Cognition, 10*, 576-590.

Lupker, S. L., & Williams, B. A. (1987). *When do rhyming primes inhibit target processing?* Paper presented at the annual meeting of the Psychonomic Society, Seattle, Washington.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word-recognition. *Memory & Cognition, 2*, 309-321.

Stemberger, J. P. (in press). Wordshape errors in language production. *Cognition*.

ACKNOWLEDGEMENTS

This research was supported by NIH Grant NS25502, awarded to Gary Dell. Correspondence should be sent to Robert R. Peterson, Psychology Department, University of Rochester, River Campus, Rochester, New York 14627. [Electronic mail: bobp@prodigal.psych.rochester.edu]