A New Model of Letter String Encoding: Simulating Right Neglect Dyslexia

Carol Whitney and Rita Sloan Berndt

Department of Neurology
University of Maryland School of Medicine
22 South Greene Street
Baltimore, Maryland 21201

cwhitney@cs.umd.edu
rberndt@umaryland.edu

Keywords: letter position, neglect dyslexia, word recognition

Introduction

Cognitive models of how oral reading is normally accomplished are typically developed and tested on the basis of experimental data gathered from skilled readers. Increasingly, however, information from other sources is playing an important role in the elaboration of such models. Detailed analyses of the effects of focal brain lesions on skills such as oral reading can place limits on the possible types of relationships that might exist among hypothesized processing components (see, e.g., Shallice, 1988). Of special interest are analyses of the relationship between the target words patients are attempting to read and the types of errors they produce, which may implicate breakdown of semantic, orthographic and/or phonologic processing. A second methodology that has begun to play an important role in the elaboration of cognitive theories of reading involves the development and testing of computational models. Implemented computer simulations of oral reading have demonstrated the feasibility of such hypothesized processing details as the distinction between serial assembly of sub-lexical units for unfamiliar words and parallel access to stored orthographic units for familiar words (Coltheart et al., 1993). Frequently, neuropsychological data are used to test computational models, which should be able to be degraded with simulated "lesions" to reproduce clinical patterns (Coltheart et al., 1993; Plaut & Shallice, 1993).

This paper describes a theoretical model of some aspects of oral reading that was developed to be consistent with experimental data from normal subjects and with existing neurobiological theory. Computer simulations based on this model can be "lesioned" to reproduce an error pattern obtained from a group of adult patients with left hemisphere lesions. The development of this model forced explicit consideration of a number of issues concerning normal word processing, especially the important issue of how letter order is encoded in words. We will describe the genesis of the model, including the motivation for modifications made during its development, to illustrate the value of computational modeling for investigating cognitive processes and their impairments.

Neglect Dyslexia and Positional Bias in Visual Errors

The error pattern at issue here is one in which patients substitute orthographically related words – words with substantial overlap of the targets' letters – for the words they are trying to read. The pattern of orthographic substitution is not random; rather, letters on the side of the target word opposite the patients' hemispheric lesions are much more likely be lost from the response than are letters from positions on the same side of the word as the lesion. A pattern of word substitution error in which responses overlap targets only for the left or right parts of words has been termed "neglect dyslexia". As the name implies, this reading pattern has been interpreted to be a manifestation of hemispatial neglect, affecting all objects appearing in the side of space contralateral to the patient's lesion. Neglect dyslexia was first described in a group of six patients with right hemisphere lesions who showed generalized left-sided neglect but no language impairment other than errors substituting letters on the left sides of words (e.g., level → novel; Kinsbourne & Warrington, 1962). Many similar cases of left neglect dyslexia have been reported for patients with right hemisphere lesions, left spatial neglect, and no language impairments (see Riddoch, 1990, for review).

The pattern of right neglect dyslexia following left hemisphere lesion has also been reported, though less frequently. Several early descriptions are available of patients who substituted words sharing only the left-most positions with the targets, and these patients did not demonstrate neglect in drawing and other clinical tasks (Casey & Ettlinger, 1960; Warrington & Zangwill, 1957). More recently, a pattern consistent with right neglect dyslexia has been described for patients with left hemisphere lesions, but the clearest cases have involved individuals who were naturally left-handed (Caramazza & Hillis, 1990a; Warrington, 1991). The most thoroughly studied of these cases showed frank right neglect and no aphasia (Caramazza & Hillis, 1990a). This patient substituted words overlapping targets only in early letter positions (e.g., journal → journey), and she did this regardless of whether words were presented horizontally, vertically, in mirror-reversed format or were spelled aloud to her. That is, the "neglect" of the ends of words was maintained even when those end letters did not fall in the right side of space. Based on these and other findings, Caramazza and Hillis argued that the patient suffered from an attentional impairment that operated over an internally generated spatial representation of the word. Further, they explicitly interpreted this finding as indicating that a spatial representation of written words is computed during the process of normal reading. (Caramazza & Hillis, 1990b).

The hypothesis that a spatial representation of written input is an element of normal word processing is a theoretically important claim. In normal readers, a spatial representation of letter input could function to encode letter order transiently, i.e., in the very short period ($<$ 40ms, Perfetti & Bell, 1991) before it begins to be supplanted by a phonetic code. Right neglect dyslexia is then interpreted as a form of "neglect" of this spatial code, and is thus necessarily linked to other manifestations of spatial neglect (see Hillis & Caramazza, 1995, for discussion).

One complication for this point of view is the occurrence of neglect-like reading errors among patients with left-hemisphere lesions and no clinical signs of spatial neglect (e.g., Warrington, 1991). Word substitutions that favor retention of early letter positions have long been described among the "visual" errors of left-hemisphere damaged patients who also show a wide range of other types of language impairments, but no reported spatial neglect (Morton & Patterson, 1980; Shallice & Warrington, 1975). In fact, retention of early letter positions appears to be a relatively common finding among the errors of patients with Deep and Phonological Dyslexia (see also Buxbaum & Coslett, 1996; Greenwald & Berndt, in press). To investigate this issue further, and to provide some actual error data that could be simulated, we carried out a retrospective analysis of an error corpus gathered from patients with focal left hemisphere lesions and consequent reading impairment.

An Analysis of Word Substitution Errors in Patients with Left-hemisphere Lesions

Patients were selected to participate in an investigation of the underlying causes of sub-lexical reading impairment. Thus, all patients demonstrated difficulty reading nonsense words that was disproportionate to their ability to read real words. All patients were premorbidly right handed adults who were skilled readers prior to suffering a left hemisphere cerebrovascular accident not less than three months prior to initiation of the study. The initial cohort of eleven patients is described in detail in Berndt et al. (1996). This group showed a range of aphasic impairments from severe non-fluency with agrammatism to a mild anomia. The ability to read real words (N=349) ranged from .18 to .97 correct, while nonword reading (N=20) ranged from .00 to .60 correct.

This group of patients produced a wide range of errors when reading words, including substitution of words semantically related to targets, perseverations of previous responses, and phonetic distortions. Of primary interest for present purposes, most of the patients with difficulty reading real words produced at least some errors in which orthographically related words were substituted for targets. These so-called "visual" errors have been defined as word substitutions that share at least half the letters of the target (Coltheart, 1980). Five patients of the group produced visual errors as a substantial proportion ($> .25$) of their errors, and these data were selected for further analysis (Berndt & Haendiges, 1997).

A letter position analysis was carried out on these visual errors (N=201) by aligning the target and response words from left to right and calculating the proportion of letters from each position that appeared in both words. For each patient, the proportion of letters retained in the response at each position was calculated across all of the visual errors, and was expressed as proportion of target letters retained per position. Several additional analyses were conducted to assure that this absolute scoring of position did not introduce scoring artifacts (Berndt & Haendiges, 1997). A strong positional effect favoring early letter positions was found for all patients (mean proportion letters retained across patients from first to last position: .80, .72, .55, .32, .27, .20). Similar positional biases were found for patients' lexicalizations of non-word targets, and a somewhat attenuated positional effect was found when all word substitution errors were scored for position (see Berndt & Haendiges, 1997).

Other aspects of the errors produced by these patients were consistent with previous descriptions of "neglect dyslexia". First, there was a strong tendency for patients' responses to be approximately the same length as the targets: there were no substitutions of "cat" for "caterpillar," for example, or vice versa. Most errors were within one letter of target length. Second, the absolute position at which letters tended not to be retained changed with increasing word length. That is, more letters were retained in the early positions of longer words than of shorter words. Another feature of these errors, which has not previously been discussed, was that only 5% involved letter transpositions. Most errors reflected deletion of letters (e.g., fact → fat) or insertions (e.g., plane → planet) or a combination of both (e.g., corn → cord, frog → frost), rather than errors that permuted the letters of the target (e.g., note → tone). These characteristics of the patient data, in addition to the positional effects, were considered in evaluating the simulations to be described below.

Although the patients who produced these errors were not tested extensively for the presence of right neglect, none of them demonstrated obvious neglect on clinical screening. In contrast, all patients showed some type of language impairment, which in several cases was quite severe. Thus, in considering possible sources for the positional effects found among this group, there was no compelling reason to assume that a degraded spatial representation of target words was a necessary component in the generation of this error pattern. Rather, it seemed appropriate to consider other ways in which letter order could be encoded during normal written word processing that could produce this pattern consequent to impairment within the language system.

Modeling and Letter Order Encoding

Our goal was thus to develop a model of written word processing that could accommodate the finding that patients with focal left hemisphere lesions tend to substitute words that overlap the first few letters of the target. More specifically, the model's structure, when degraded, should fail in a manner indicating that lexical access is based on a partial representation that includes information about the identity and order of words' initial letters and about target word length. The assumption is that the disorder responsible for this pattern (similar to "right neglect dyslexia") is substantively distinct from the disorder responsible for left neglect dyslexia, which seems much more clearly linked to an impairment involving all information processing in left extrapersonal space (see also Ellis et al., 1993). Rather than demonstrating generalized problems with information presented in right space, patients with the pattern of right neglect dyslexia tend to show problems only with language. Specifically, the patients whose data are presented here showed only one common symptom across the entire group: impaired ability to "sound out" unfamiliar letter strings. It is unclear at this time whether or not this symptom is functionally related to the word substitution pattern favoring early letter retention. This possibility is discussed further below.

In light of the clear language impairments demonstrated by our patient group, and the lack of clear attentional/spatial impairments, the architecture of our model was designed to be consistent with findings in the cognitive literature about normal written word processing. As reviewed below, the available data are consistent with the idea that the relative order of letters in a string is an important factor in normal subjects' processing of written letter strings. These studies do not support the hypothesis that the representation of letter identity and order is spatial in nature, but it is less clear precisely what sort of representation they do support. The complex pattern of results related to normal letter string encoding becomes especially evident when attempting to develop an explicit computational account of the processes involved.

Previously published computational models of word recognition and pronunciation have adopted one of two mechanisms for the encoding of letter order, neither of which is entirely satisfactory. Several connectionist models of word recognition (McClelland & Rumelhart, 1981) and pronunciation (Coltheart et al, 1993; Whitney, Berndt & Reggia, 1996) have employed a "channel-specific" coding scheme that requires the representation of every letter in each position. Presentation of a letter string then activates only constituent letters in the correct position. This straightforward approach to letter order assures that later processing will be based on the correct order of letters, but it requires a high degree of item redundancy. Morever, channel-specific coding cannot account for findings of relative position priming, where facilitation of word targets occurred when the relative order, not the absolute position, of letters was preserved in the prime (Humphreys et al., 1990; Peressotti & Grainger, in press). Nor can it account for position indepedent priming where facilitation of alphabetic identification of trigrams occurred when both absolute position and order were violated in the primes (Peressoti & Grainger, 1995).

Another approach to encoding information about order in computational models of reading is to postulate the existence of multiletter clusters, usually trigrams, as a functional unit (Mozer, 1987; Mozer & Behrmann, 1992; Seidenberg & McClelland, 1989). The use of cluster representations produces a model that is sensitive to the local context in which letters appear rather than to their specific position in the string. Consequently, cluster coding of letter input can better account for relative position priming results than can models with channel-specific representations. However, the postulation of letter clusters begs the question of how these sublexical, supraletter clusters are generated during orthographic processing. It is difficult to see how trigrams could be activated directly from orthography without any recognition of the constituent letters, and such a process is inconsistent with evidence for facilitatory priming effects on individual letters (Grainger & Jacobs, 1991). However, if contextual units are activated from representations of the constituent letters, then the question of how letter order is initially represented remains unaddressed.

Based on these considerations, our goal in modeling the patient data was to construct a theoretical framework for written word recognition that was consistent with what is known about normal processing, and to implement that framework in a model that could be lesioned to simulate the errors produced by the patient group. At the same time, we attempted to address some of these outstanding issues concerning the representation of letter order in a manner that accommodates current neurobiological thinking. The remainder of the paper is divided into four sections. First, we present our initial formulation of order encoding and its underlying rationale, including a discussion of its limitations. Next, we describe an elaborated version of the model which addresses those limitations. Then we present the results of three simulations based on this model. Finally, we propose a potential source of the hemispheric differences that produce the symptoms of right and left neglect dyslexia.

Letter Order as an Activation Gradient

The central tenet of our conception of the representation of letter strings is that individual letters within the string (from first to last) are not all activated to the same degree when a letter string is presented. Rather, we conceive of an activation gradient across letter nodes, with highest activation of the first letter and gradual fall-off to the end position. For example, the string "CAT" would be represented by clamping the activation of C to a maximal value, A to a value less than C, and T to a value less than A. Under this scheme, there would be a single set of letter nodes, and the position of a letter within a string would be represented by its level of activation.

This general framework has some support from studies of normal string processing, especially for the idea that initial letters in a string are the most highly activated. For tachistoscopically presented strings, probability of letter recall and identification falls off from left to right (Lefton et al., 1978; Hammond & Green, 1982; Mason, 1982). Masked form priming studies have shown the initial letter to be the most facilitatory in target identification. For nonword primes having one letter in common with word targets, target naming was facilitated only when the common letter was in the initial position (Humphreys et al., 1990). For a task involving word primes and identification of target letters, identification was facilitated only when the target letter occurred in the initial position of the prime (Grainger & Jacobs, 1991).

Other studies have indicated that strings differing in their initial letters are more dissimilar than strings differing in medial positions. For orthographically legal nonwords constructed by transposing two adjacent letters of a word, lexical decision latencies were slowed when the transposed letters were in the middle of the word, but not when they were at the beginning (Chambers, 1979). Interference in a word identification task occurred when a prime differed from the target in medial positions, but not in initial positions (Perea, 1998). In string comparison studies, subjects took longer to judge that two strings were different when the mismatch occurred at the end of the string than when the mismatch was at the beginning (Ratcliff, 1981; Proctor & Healy, 1985, 1987). Taken together, these results indicate that skilled readers show preferential encoding, or enhanced processing, of letters in the initial position. As noted below, the literature on normal letter string processing contains considerable evidence that positional encoding of letters is more complex across the entire string than these findings would suggest. However, our initial attempts to simulate these data focused on the feasibility of capturing a simple left-to-right processing advantage without using either channel-specific or contextual letter units.

Activation Gradient Simulations (I)

In our pilot simulations, we investigated the feasibility of representing letter positions via an activation gradient. Here we briefly summarize the structure and results of our initial simulations. The network consisted of two layers, a set of letter nodes (the input layer) and a set of word nodes (the output layer). The letter level was comprised of a node for each single letter (A-Z), and a node for each double letter (AA-ZZ). The word level consisted of a node for each monosyllabic word from the NETTalk corpus (Sejnowski & Rosenberg, 1987), yielding approximately 3,500 word nodes.

To represent an input string, the activation of the corresponding letter node was set to 0.7^{pos-1}, where *pos* denotes position within the string. For example, the input string "TREE" was represented by setting T to 1.0, R to 0.7, and EE to 0.49. If a letter was repeated, but not as a double letter, the activation of that letter node was set to the sum of the activations for each position.

The weights on the letter-to-word connections were set, rather than learned. The weight vector for each word was set to the letter activation vector corresponding to that word, where the length of that vector was normalized to 1.0. The activation of a word node was the dot product of its weight vector and the letter activation vector. The word node with the highest activation constituted the output of the network. To lesion the network, noise was added to the activation of each word node. As discussed below, the selection of noise at the word level as a "lesion" mechanism was intended to represent degraded input from impaired sublexical translation mechanisms.

The pilot simulations were encouraging. Each word in the database was correctly recognized. When lesioned, the simulation produced errors that were positionally biased from left to right, and preserved target length. Thus, we were able to implement a novel string representation without a channel-specific representation of letter order, and without introducing context units. Further, a "word-centered" spatial representation was not required to reproduce salient aspects of the patients' errors.

Limitations of the Pilot Simulations

Although this pilot simulation provided general support for our approach, several aspects of the data suggested the need for modification. First, the simulation of the patient data failed to reproduce some elements that we interpreted to be important characteristics of patient performance. For example, the patients' errors consisted of substitions, omissions, and/or additions of letters, with very few transpositions. However, a majority of the error responses from the pilot simulations involved transpositions. In addition, the shape of the simulated positional effect was different from the patients' with respect to the initial positions. The patients showed similar retention levels across the first and second positions. However, the retention level at the second position was significantly lower than at the first position in the simulations. A more general problem was that the graded activation representation was not very robust. Activation levels of anagrams (e.g., CLAPS and CLASP) were very similar to each other, and even a small amount of noise caused errors.

In addition to these problems with the patient simulations, we were also concerned that the pilot model was unable to accommodate some complex results from the normal literature concerning the status of terminal letters. Because the final letter was the least likely to be retained in the patient data, we assumed that the lowest level of activation occurred in the final position. A low level of activation at the terminal position is also consistent with results of string comparison reaction times reported by Proctor and Healy (1987). In one set of experiments (the "order" task) normal subjects were to compare two strings of four consonants and report "same" if both strings had the same letters in the same order, and "different" if any change was detected. Latencies to make correct "different" responses showed significant effects of the position of the difference, regardless of whether the difference involved a replacement of one of the consonants or a shift of position. Changes were easier to detect at the beginnings compared to the ends of strings. However, in another experiment (the "item" task), subjects were to respond "same" if both strings contained the same letters, regardless of their order. Under these conditions in which subjects were explicitly NOT to attend to order, subjects were faster to detect a changed item at the end of the string than in the second position. Thus, processing of the terminal letter in the string was not always disadvantaged relative to earlier positions, but changed in these experiments as a function of task.

There was also conflicting evidence reported about the activation of the final letter in the masked form priming studies of Humphreys and co-workers (1990). For nonword primes having two letters in common with word targets, target naming was most facilitated when the common letters were in the initial and final positions. Yet an analysis of errors made under the neutral prime condition (no letters in common with the target) showed that the target's final letter was the least likely to be correctly identified. In other studies, the final letter was also the least likely to be correctly identified in briefly presented trigrams (Hellige et al., 1995) and words (Montant et al., 1998). In an alphabetic indentification task using trigrams, false positives were the most likely to occur when the foil appeared in the third position, indicating that the final position had the lowest level of activation (Peressoti & Grainger, 1995).

In contrast to these findings of low levels of activation for final letters, several other studies have indicated that the final letter has an increased activation relative to the internal letters. For tachistoscopically presented strings, efficacy of recall and target identification increased in the final position (Lefton et al., 1978; Hammond & Green, 1982; Mason, 1982). Interference in a word identification task occurred when a prime differed from the target in medial positions, but not when it differed in the final position (Perea, 1998).

It is difficult to reconcile these conflicting results. In all studies, the initial letter appeared to be the most active, with activation declining from left to right across the string. Yet evidence for the activation of the final letter was contradictory: some studies indicated that it was the least active of all the letters; others indicated that it was more highly activated than the internal letters. Obviously, our initial activation gradient model could not accommodate these findings.

A Theoretical Framework for Letter String Encoding

Although the simulations using an activation gradient across letter units reproduced some of the positional effects in the patient data, it appeared that an additional constraint on relative letter order would be required to decrease the number of transposition errors and to provide a more robust representation. Such a constraint could be achieved by the interpolation of letter cluster units, specified for order, between letter and word representations. As noted above, however, the use of cluster units in computational models of letter string processing has been relatively unmotivated. We wanted to develop a framework within which all of the hypothetical levels we proposed were based on evidence that they play a role in normal processing. Further, we thought it necessary to propose some mechanism by which such contextual units become activated. Accordingly, in this section we set out the theoretical framework for a model of letter string encoding that is motivated by findings in cognitive psychology and neurobiology. This theoretical model serves as a basis for computer simulations, which are descibed in the following section.

In trying to understand the nature of letter activations and their possible relationship to contextual units, we conceived of a novel model of letter position coding. We postulate that a letter string is encoded by a temporal firing pattern across letter nodes: position is represented by the precise timing of firing relative to the other letters. The induction of this firing pattern results in an activation gradient. Thus, the activation gradient does not in itself represent position, but rather comes about as a by-product of the representational process responsible for letter node activation. Such a model provides the flexibility to account for the contradictory findings, noted above, with regard to final-letter activation.

Contextual units are activated by the sequential dynamics of the letter nodes, and provide a mechanism for decoding this temporal representation. We chose bigrams as the contextual unit since they are the smallest possible such unit. In addition, there is evidence from studies of normal letter string processing that bigrams indeed comprise a functional processing unit (Whitely & Walker, 1994, 1997), and arguments have even been made concerning the source of these units. For example, Whitely and Walker (1994, p. 469) suggest that bigrams may emerge as units by virtue of their frequent correspondence to higher level representations, including syllable elements and familiar morphemes. Our approach to the genesis of bigram representations is different from that of Whitely and Walker's in that we postulate that bigrams are activated from processes originating at the letter level.

In this theoretical model, we will use the word "node" to refer to the basic computational unit, which may be thought of as corresponding to a neuronal assembly.

Functionally, a node recognizes the occurrence of a symbol. Following Hopfield (1996), we assume that a pool of nodes recognizes each symbol, and that a subset of that pool responds to each occurrence of a symbol. If the same symbol appears more than once within an input string, a different subset becomes active for each occurrence.

The model consists of three levels of nodes (letters, bigrams and words), each of which has unique activation dynamics. At the letter level, nodes are specialized for creating a sequential firing pattern from graded inputs. At the bigram level, nodes are specialized to recognize temporally ordered pairs of letters. Bigram nodes convert the temporal encoding from the letter level to a non-temporal representation. At the word level, nodes are specialized to recognize combinations of bigrams. For example, the string CART would be represented in the following way: at the letter level, C would fire, then A, then R, then T; at the bigram level, CA, AR, RT, CR, AT, and CT would become activated; at the word level, CART would be the most active of the word nodes. This structure is illustrated in Figure 1. The following sections address each level in greater detail.

[Figure 1 about here.]

Letter Level

The idea that letter position within a string is represented by the precise timing of neural spikes is consistent with current neurobiological models of information encoding. Hopfield (1995) has proposed that quantities are represented by the explicit timing of action potentials, rather than by their "firing rate". In this "phase-advance" model, encoding neurons undergo internal, sub-threshhold oscillations of membrane potential. The magnitude of an input to such neurons determines when threshhold is exceeded. For a small input, threshhold is not exceeded until late in the cycle when the cell's oscillation brings its potential near threshhold. For a larger input, threshhold is exceeded earlier in the cycle. Thus, the size of an input is represented by spike timing relative to the oscillatory cycle. This scheme implies that individual spikes are much more important than has traditionally been assumed. Indeed, recent studies have shown that single spikes encode signficant amounts of information (Reike et al., 1997), and that spike timing is precise and reproducible at a millisecond time scale (Victor and Purpura, 1996; Berry et al., 1997; de Ruyter van Steveninck, 1997).

Our intention is not to incorporate specific neurobiological details into our formulation of letter string encoding, but to provide an account that could ultimately be linked to such details. The framework offered here attempts to provide an interface between cognitive data concerning positional encoding and current thinking about underlying neural activity. Hopfield (1995) has proposed the phase-advance model as a general method of information encoding used by systems performing non-serial processing (e.g., in the olfactory system), as well as serial processing.

Given that all other modalities of language input and output (hearing, speaking and writing) involve serial processes, it is not unreasonable to assume that a spatial letter string is converted into a temporal representation at some point. This temporal representation does not arise from a serial input process. Rather, parallel input is converted to a temporal representation at the letter level of our model. This is accomplished by graded inputs in conjunction with internal subthreshhold oscillations and lateral inhibition. The graded inputs are such that the first letter receives the most input, the second letter receives the next highest amount, and so on. The letter receiving the highest level of input fires first because it reaches threshhold before the others. The letter receiving the second highest level of input fires next, etc. Suitable input levels and lateral inhibition assure that only one letter fires at a time. Our model assumes the existence of graded inputs; possibilities regarding their source are discussed below.

More precisely, we assume that all letter nodes undergo the same oscillations, $C(t)$ with period g, of internal state, and that all letter nodes have the same firing threshhold, T. We will consider activity during a single oscillatory cycle $0 < t < g$, where $t = 0$ is a time for which C is at its minumum. Let L_i denote a letter node, and V_i denote its internal state. Input to L_i is measured in terms of its effect on V_i. Let E_i denote excitatory input to L_i. For simplicity, we take E to be independent of t. Prior to firing, a node's internal state is given by $V_i(t) = E_i + C(t) - I_i(t) + v$, where $I_i(t)$ denotes inhibitory input and v is a constant. When $V_i(t) > T$, L_i fires and sends inhibition to all other letter nodes. Immediately after firing, V_i falls to v due to internal post-firing inhibition. The rate of rebound of V_i depends on E_i. We assume that lateral inhibition and post-firing inhibition interact in a multiplicative manner. Lateral inhibition has a small effect if the receiving node has not fired recently, and strong effect if has (sufficient to inhibit the node for the rest of the cycle). Thus a node can fire repeatedly until it receives lateral inhibition.

Given this framework, a string of letters, $l_1 l_2 .. l_n$, can be represented by the sequential firing of letter units. Here we assume that n is small enough so that all letters can be represented within a single cycle. Based on experimental evidence to be discussed below, we assume that $g \approx 200$ms, and that a maximum of about 8 letters can be encoded within a cycle. Let s_i denote the first time that L_i fires, e_i denote the last time, and k_i denote the total number of times that L_i fires (activation). Denoting the node that recognizes l_i as L_i, it should be the case that $s_1 > s_2 .. > s_n$, and $e_i < s_{i+1}$. This constraint can be satisfied by having $E_i > E_{i+1}$ and $V_i(s_{i-1}) < T$. This assures that the letters fire one at a time, in order.

We assume that internal letters ($i > 2$ & $i < n$) attain the same minimal level of activation. In accordance with the experimental evidence for the importance of initial letters, we assume that the large inputs to L_1 and L_2 cause them to become more active than this minimal level. So $k_1 > k_2 > k_i$, for $i > 2$ & $i < n$. Since the final letter does not receive lateral inhibition after it commences firing, it will continue to fire until it can no longer exceed threshhold due to the decrease of C. Thus, k_n depends on s_n. For small enough n, $k_n > k_{n-1}$. Thus, although the final letter receives the lowest level of input, its activation may be higher than that of the internal letters.

So at the letter level, parallel inputs are converted to serial outputs; simultaneous graded inputs, in conjuction with the dynamics of the letter nodes, create a temporal firing pattern. This firing pattern encodes the order of the letters within the string. As a side effect of this process, the initial letters and the final letter have a higher activation level than the internal letters. Thus, the input gradient to the letter nodes creates an activation gradient across the letter nodes. However the activation gradient does not necessarily have the same shape as the input gradient because the activation gradient results from the interaction of the input gradient with other factors (namely, the internal states of the letter nodes, and lateral inhibition between letter nodes). The input gradient is strictly decreasing from initial to final position, while that is not necessarily the case for the activation gradient. This principle can account for the conflicting evidence regarding the apparent activation of the final letter in a string. Simulations to illustrate this effect are presented below.

Bigram and Word levels

The postulation of a temporal representation of letter order poses the problem of how such a representation might be decoded by the word nodes. We see two possibilities: either word nodes directly decode the temporal patterns, or intermediate units recognize temporal sub-patterns, and word units decode combinations of these non-temporal units. Corresponding to the first option, Hopfield (1995) proposes that recognition of a temporal pattern involves learning time delays (physical connection lengths) so that all inputs arrive at the correct decoding unit simultaneously; the target unit then performs coincidence detection.

Alternatively, short sequences could be recognized directly. That is, specialized nodes fire when two inputs occur in the proper order. The bigram level is based on this assumption. A bigram node becomes active if it receives suitable inputs within the time span of an oscillatory cycle. We hypothesize that neuronal assemblies exist which fire only if input A is followed by input B. Such an assembly would not fire if only A were received, or if B were received prior to A.

A bigram node becomes activated even if its corresponding letters were not immediately contiguous in the string. Let B_{ij} denote a bigram node that becomes activated by input from L_i followed by input from L_j. We assume that the activation level of B_{ij} is proportional to the product of k_i and k_j, and decreases with the time interval between between e_i and s_j. The slope of this decrease is assumed to be steep initially, and then to flatten out. This means that other things being equal, bigrams representing contiguous letters are more active than bigrams representing non-contiguous letters, while the activations of bigrams representing letters separated by one or more letters are comparable. This proposed scheme is much more robust than a time-delay architecture. Redundant information is encoded across the bigram units, whereas the time-delay scheme depends on learning precise connection lengths.

At the bigram/word interface, our model follows the classical firing-rate model. That is, the activation of a word node is a function of the dot-product of its weight vector and input vector. The input vector is comprised of the activations of all bigram nodes subsequent to representing the input string across the letter nodes. The weight vector is set to the vector comprised of the bigram activations corresponding to the word; i.e., the connection weight between a bigram node and word node is set to the bigram node's activation subsequent to representing the word across the letter nodes.

The Model and Experimental Data

Specific elements of the theoretical model proposed above are supported by data from a variety of sources. In this section we review some additional data from the literature that constrains our formulation of particular elements of how letter string information is encoded. The actual dynamics of these elements are further supported by the results of three computer simulations. Simulations at the letter level illustrate the induction of a temporal firing pattern from graded inputs, and show how the model can account for the conflicting data on the activation of the final letter. String comparison simulations show how the proposed representation at the letter level can account for the complex positional results from Proctor and Healy's (1987) order task. Finally, returning to the data that launched this investigation, we present simulations of the patients' errors.

Letter Level

It has been proposed that oscillatory activity in the brain near 40 Hz (gamma frequencies) is related to cognitive processing (Tiitinen et al., 1993). There is evidence that individual 40 Hz waves are related to individual auditory stimuli (Joliot et al., 1994). It has been suggested that short-term memories are encoded on 40 Hz sub-cycles of a low-frequency (5 to 12 Hz) oscillation (Lisman & Idiart, 1995). This proposal is consistent with the observation of nested oscillations recorded in the human cortex in response to auditory stimuli (Llinas & Ribary, 1993).

We suggest that such oscillations also underlie visual language processing. In our model, each letter position corresponds to a successive 40 Hz sub-cycle within the oscillatory period $g \approx 200$ms. This proposal is consistent with some curious results from a study involving sequential letter presentation. In this study, the letters of 8-letter words were presented one at a time across a horizontal row (Mewhort & Beal, 1977). The interval between successive letters (ISI) was varied, and performance was measured as probability of correctly identifying the word. For ISI's of 0ms, 50ms, and 125ms, performance was 98%, 70%, and 50%, respectively. However, for an ISI of 250ms, performance rebounded to 65%, rather than continuing to fall off. Our interpretation is that the sequential presentation interfered with the normal phasic coding of letter position. Letter presentations were maximally out of phase with respect to normal at an ISI of 125ms. Performance levels for 50ms and 250ms were similar, consistent with a cycle length of 200ms. A variation of this task, where the unit of presentation was a syllable rather than a letter, provides further evidence of periodicity. Errors rates increased as ISI increased from 0 to 250ms, decreased back to the 0ms level with further increases of ISI from 250 to 500ms, and increased again from 500 to 625ms.

Temporal Encoding Simulations

We have performed simulations to illustrate the proposed firing pattern at the letter level. The simulations were implemented by choosing specific values for the functions described above in the the specification of the letter level. The details are given in Appendix A.

The activation pattern of the nodes relative to each each other can be manipulated by varying the overall level of input (the absolute size of the $E's$). For a high input level, k_n (the activation of the node representing the final letter) may be greater than k_{n-1} (the activation of the node representing the penultimate letter), as described above. For a low input level, the final node may not fire at all or may fire only a minimal amount, as s_n (the onset of firing for the node representing the final letter) is pushed near to $g/2$ (the time at which the internal oscillation is at its maximum). This effect is illustrated in Figure 2.

[Figure 2 about here.]

We propose that variation in overall input level is one component contributing to the contradictory evidence regarding the importance of the final letter in the normal literature, as discussed above. To review, although some studies showed increased activation of the final letter relative to string-internal positions, there were five experimental results in which the final letter seemed to be the least activated of all letter units. These findings were generated by a number of distinct experimental paradigms, including the neutral-condition masked form priming studies of Humphreys et al. (1990), central fixation trigram identification results from Hellige et al. (1995), word identification data from Montant et al. (1998), false positive responses in Peresotti and Grainger (1995), and the “order” condition of Proctor and Healy (1987). The activation dynamics postulated in our model can accommodate low levels of activation for final letters under these experimental conditions.

We assume that the final letter is normally more active than internal letters due to lack of lateral inhibition subsequent to firing. However, when overall input level is reduced, the final letter is the most vulnerable. When stimulus presentation is very short, input level is reduced; the final letter becomes the least activated. Indeed, it is the case that in Humphreys et al. (1990), Hellige et al. (1995) and Montant et al. (1998) target presentation durations were very short, approximately 40ms. In contrast, all experiments that showed a privileged status for the final letter used target durations of 75ms or greater. Thus, the data are consistent with the idea that the final letter fires until it can no longer reach threshhold due to the downward phase of the oscillatory cycle. At very short presentation durations and low levels of input, the final letter fires very late in the cycle, and becomes the least active of all the letters. For longer presentations and higher levels of input, the final letter starts to fire earlier in the cycle, and fires more than the internal letters.

However, there were contradictory data from within the Humphreys et al. (1990) study; the authors found the initial/final combination to be the most facilitatory for priming of all possible two-letter combinations. How could this combination be the most facilitatory when the final letter is the least active of the letters? The key is that priming measures the after-effects of firing, not the level of firing itself. We have proposed that letters become inhibited if they receive lateral inhibition after firing. Interpreted this way, the initial/final combination is the most facilitatory because it does not involve the internal letters, which have become inhibited as a result of prime presentation. Although the initial letter has also received inhibition, its internal state is still above normal due to its high level of input. This idea is consistent with a study showing that nonword primes have inhibitory effects on letter identification within a string for medially located targets, but not for initially or terminally located targets (Grainger & Jacobs, 1991). However, in the Humphreys et al. (1990) study, the two-letter combination consisting of the internal letters was found to be facilitatory compared to the neutral condition. If, as proposed, the internal letters are inhibited after prime presentation, how does this facilitation occur? We propose that it occurs at the next level of processing, the bigram level, because activation of the internal bigram of the target is facilitated.

In the other studies which implied low levels of activation for the final letter (Peressoti & Grainger, 1995; Proctor & Healy, 1990), targets were displayed until a response was given by the subject. Thus, short presentation durations cannot account for these results. Despite the extended presentation duration, the Peressotti and Grainger (1995) result can be understood as a reflection of the low level of input to the final position. In that study, 3-character strings were presented, and subjects were to determine whether or not the string consisted entirely of alphabetic characters. Foils consisted of two letters and one non-letter (e.g., A%B). When the non-letter appeared in the final position, false positives were the most likely to occur, indicating that the final character was the least well represented. We propose that a non-letter final character does not display the usual advantage that results from lack of lateral inhibition because its activation dynamics are different from that of letters. Letters are interconnected in a network that processes strings, while non-letters are not. Thus, the low level of input to the final position is shown directly when that character is not a letter. This proposal is consistent with studies showing that strings of digits and letters are processed differently from strings of other characters. Digits, like letters, have shorter detection latencies at the ends of a string. However, strings of other types of symbols do not show this pattern; latencies are increased at the ends (Hammond & Green, 1982; Mason, 1982).

Our account of the Proctor and Healy (1990) result involves a higher level of processing than the letter level, namely that of string comparison. The fact that our account depends on the comparison process is consistent with the finding that the role of the final letter varied with the type of comparison task ("order" versus "item") in those experiments. In the "order" task subjects were to determine if two strings were comprised of the same letters in the same order. Two types of differences between strings were analyzed. For "replacement" pairs, the two strings differed by the substitution of a letter at a single position, i.e., ABCD and ABXD. Response time increased monotonically with replacement position, i.e., XBCD was the fastest to be rejected, and ABCX was the slowest. Thus, no privileged status for the final letter was observed. For "permutation" pairs, the two strings were comprised of the same letters in different orders, i.e., ABCD and DBCA. Response time was a complex function of total displacement of constituent letters, and of positional matches.

We propose that the final letter was indeed more highly activated than the internal letters, but that this was not evident in the pattern of response times for replacement pairs due to the nature of the comparison operation between the two strings. We have implemented a simulation that illustrates this hypothesis. Moreover, the simulation also captures the pattern of response times observed for the permutation pairs.

String Comparison Simulations

The simulation is based on the premise that letter position is represented by the timing of firing relative to an underlying oscillatory cycle. Thus, representation of position can be considered to be a waveform over time. We model comparisons between strings as an interaction between the constituent waveforms over the entire oscillatory period. Using a trial and error process, we arrived at parameters that yielded a good fit to the experimental data. The underlying period is represented by time slots 1-120. Firing is represented as occurring over a block of time, rather than as individual spikes. A node's activation corresponds to duration of firing. The first letter in a string fires during the first 50 time slots, the second letter fires during the next 25 slots, the third letter during the next 13 slots, and the fourth and final letter fires for the next 25 slots. Specifications of the waveforms, and the comparison function are given in Appendix B.

Figure 3 shows the results of our simulations compared with the experimental data. Note that the response times for the replacement pairs are monotonically increasing despite the fact that the fourth letter is active longer than the third letter (25 slots vs. 13 slots). This result occurs because the comparison function depends on the entire waveform, not just the activation period. Note also that the simulations reproduce the pattern of response times from the permutation pairs. The two simulations were run with the same parameters. Thus, simulations based on the premise that string comparisons can be modeled as the interaction of temporal waveforms yielded a good fit to these complex data.

[Figure 3 about here.]

Bigram and Word levels

The simulations presented above lend support to our proposed mechanisms for encoding letter order in a manner that produces an activation gradient across letter units, which is then propagated to the bigram units. Next we turned our attention to the dynamics of bigram-to-word activation in an attempt to generate the correct words from letter string input and to simulate the patient data.

Activation Gradient Simulations (II)

These simulations were similar in structure to the pilot activation gradient simulations, except that the input layer consisted of bigram nodes, rather than letter nodes. The network thus consisted of two layers of nodes, the bigram (input) layer, and the word (output) layer. The bigram layer consisted of all possible combinations of two letters, and the word layer was the same as in the pilot simulations (approximately 3,500 monosyllabic words).

Activations of bigram nodes were set in accordance with the assumptions outlined above, namely that their activation is proportional to the activations of their constituent letters, and that activation is reduced with the separation of their constituent letters. Letter activations were not modeled directly, but rather were assumed to vary with position (as simulated above). This letter activation gradient was incorporated in the specification of the activation of bigram nodes. Thus, the activation of a bigram node depended on the the positions of its constituent letters within the word being represented across the input layer. We assumed that despite the separation between the initial letter and the final letter, their high levels of activation would combine to make bigram B_{1n} highly activated. The details of the bigram activations are given in Appendix C.

As in the pilot simulations, weights between the two levels were set, not learned. The weight vector for each word node was set to the bigram activation vector corresponding to that word. The activation of a word node was calculated as the dot product of its weight vector and the input vector. The word node with the highest activation was selected as the output of the simulation. In a test of intact performance, the network correctly identified each word in the database.

Lesion Simulations

The performance of the intact model suggests that the structure we have proposed is generally feasible as a model of written word activation. However, the primary goal of the undertaking was to produce a model that could be degraded to reproduce several aspects of our patient data. Ideally, the lesion simulations could provide us with a means of testing our ideas about the processing limitations that led to the reported symptoms. As noted above, all of the patients in our sample showed some difficulty reading aloud following their stroke. Although they varied widely in their ability to read words, all of them showed difficulty sounding out non-words that was disproportionate to their ability to read words. Although our model does not incorporate a means for allowing sublexically generated phonetic input to affect activation of the word nodes, there is evidence that such input is available to normal readers very early in the course of written word processing (Perfetti & Bell, 1991; Ferrand & Grainger, 1993). For our patients, however, any information generated by sublexical print-to-sound translation would be expected to be abnormal. Based on this assumption, we "lesioned" the model by adding normally-distributed noise (representing input from impaired sublexical translation) to the word nodes in the patient simulations.

A second factor we considered in lesioning the model concerned the apparently low level of activation found for the final letter in the positional analysis of the patient data. As discussed above, we hypothesized that the last letter in a string is normally more highly activated than internal letters (because it does not receive lateral inhibition after firing), except in conditions of degraded input. We assumed that some degree of difficulty with the activation of letters might reproduce the effect of degraded input conditions and lead to low levels of activation for final letters. Although we have no independent evidence that these patients actually showed letter activation impairments (because we did not look for them at the time the patients were tested), it is possible that they did. Arguments have been made that subtle deficits in letter activation can be uncovered through the use of specially designed tasks, even in patients who do not show clinical signs of such impairments (e.g., Reuter-Lorenz & Brunn, 1990). Based on this idea, the activation of B_{1n} was not adjusted to a high level in the lesion simulations. (Note, however, that the connection weights remained unchanged; the weights on such bigram-to-word connections retained their original high values.) This reasoning suggests that it should be possible to further test the relationship between letter activation impairments and retention of the final letter in word substitution errors through careful study of individual patients. If the assumptions underlying this lesion locus in our simulations are correct, the two symptoms should be correlated.

Various parameters for the distribution of noise and the positional activation values were tried in order to find those that gave the best fit to the patients' errors. The following results are based on simulations run with noise having a mean of 0, and a standard deviation of 0.6. The resulting errors were very similar in form to those made by the patients. Errors consisted of insertions and deletions, with very few transpositions. Figure 4 shows the results of the positional analysis of the experimental and simulated results. This analysis was performed in the same manner as described above for scoring the patient data, where the proportion of errors retaining the letter in each position was tallied. Note that the positional curve for the simulated results captures the overall shape of the experimental curve. Retention levels fall off slowly between positions 1 and 2, and then more sharply thereafter.

[Figure 4 about here]

As discussed above, the degree of retention by position varied with target length. Figure 5 displays retention level at position 3 for each target length. It is evident that the third letter is less likely to be retained in three-letter words than in longer words. The simulations reproduce this effect.

[Figure 5 about here.]

It was also the case that errors tended to preserve the length of the target. These data are displayed in Figure 6. For each target length, the average length of all the error responses was calculated. This average was close to the target length. The simulated results also show increasing response length with increasing target length, although not as strongly. This effect was very robust and held over all parameters that were tried. We think the effect is not as strong as was observed in the patients' results due to the fact that the simulation's word database consisted of monosyllabic words only. Thus the simulation did not have access to longer words, resulting in shorter responses for targets of length 5 and 6 than the patients gave.

[Figure 6 about here.]

The simulations appear to have captured the primary characteristics of the patient data through the introduction of two "lesions" in the model's activation dynamics. Although we attempted to motivate our degradation of the model's functioning in a principled manner, further research is needed before we can assume that these simulated "lesions" represent functional cognitive deficits. Specifically, we need to investigate the ability of the model to reproduce deviations from the modal pattern of patient performance.

Discussion

We have presented a theoretical framework in which the order of letters within letter strings is encoded through the activation of bigram units from letter units. Unlike previous models employing multi-letter contextual units to encode letter order, we have attempted to account for the activation of bigram units as a consequence of the dynamics of letter activation. Letter units are activated as a temporal firing pattern across positions; this pattern is initiated by graded stimulus inputs and is modulated by lateral inhibition from subsequent letters' activations. Elements of this scheme have been shown to accommodate a number of complex results regarding normal readers' encoding of letter positions across a variety of tasks. The representations and processing dynamics postulated here were shown to result in successful activation of correct word nodes for monosyllabic words, and most importantly, could be degraded in a principled fashion to simulate a number of characteristics of the reading errors of patients with left hemisphere cortical lesions.

Our simulation of the "right neglect dyslexia" pattern of errors was accomplished without the postulation of a spatial representation of the letter strings, and without assuming that the patient data reflect "neglect" of this spatial representation. Rather, the positional results (and the other aspects of the data that were simulated) are interpreted as a reflection of breakdown within processes that are specific to language, rather than to a form of attentional impairment. One critical issue remains to be discussed.

In our model, the induction of a temporal firing pattern across letter nodes relies on the postulation of a gradient across the letter input: initial letter positions provide the strongest input, second positions somewhat less input, etc. What is the source of these graded inputs? There is evidence that graded inputs may be related to the way information is processed, or attention is allocated, by each of the cerebral hemispheres.

Many different characterizations have been given to the different "modes of processing" that describe the manner in which each of the two hemispheres deals with information (Heilman, 1995; Hellige, 1993). One such characterization from studies of normal letter string identification is that the left hemisphere processes information in parallel, with attention deployed across all letter positions relatively equally, while the right hemisphere uses a less efficient serial mechanism to process letters (Hellige et al., 1995; Reuter-Lorenz & Baynes, 1992). For example, normal subjects identified CVC letter strings better when they were presented tachistoscopically to the right visual field(RVF)/left hemisphere(LH) than when they were presented to the left visual field (LVF)/right hemisphere(RH)(Eng & Hellige, 1994; Hellige et al., 1995).

This result supports the general superiority of the left hemisphere for the processing of linguistic stimuli. However, an important (and counterintuitive) result involved the error patterns produced for each letter position as a function of the visual field to which strings were presented. With LVF/RH presentation, subjects made many more errors that involved deletion of the last letter, relative to the first letter, of the string. With RVF/LH presentation, this finding was greatly attenuated: there were more errors on the first letter, and fewer errors on the last letter, resulting in a more even distribution of errors across the string. Interestingly, when stimuli were presented to both hemispheres in a central fixation (Hellige et al., 1995), or were presented simultaneously to both visual fields (Eng & Hellige, 1994), the error pattern that emerged was the same as the right hemisphere pattern. This finding suggests that, even though the LH was the most effective at performing the task, the RH's mode of processing dominated when the stimuli were presented to both hemispheres simultaneously.

The effect produced is that with central fixation, stimuli are processed in a manner that results in superior processing of the first letter, declining across positions. Although this advantage for early letter positions appears to apply to normal (central fixation) reading, at least for languages that are read from left-to-right (Chokron & Imbert, 1993; Chokron & DeAgostini, 1995), its precise source is not entirely clear. In addition to the interpretation offered by Hellige et al. (1995), the results could reflect the manner in which the right hemisphere allocates attention across letter positions. Whatever its source, in our model the effect of these graded inputs is carried forward (for bigram and word units) such that even as the processing levels become increasingly abstract (i.e., farther removed from the visual input), the advantage for early letter positions is maintained. As shown in our initial simulations using an activation gradient across letter nodes, the entire constellation of results cannot be successfully reproduced through the simple postulation of an activation gradient. The model provides details about the effects such graded inputs could have on subsequent levels of processing.

Our model, and the "lesion" simulations that were carried out, offer an alternative to the idea that the pattern of errors found in "right neglect dyslexia" indicates an attentional deficit across a spatial representation of letter strings. The situation that emerges when the right hemisphere is damaged, and the initial positions in letter strings are subject to disruption when reading (left neglect dyslexia), clearly requires a qualitatively different explanation. In the case of right hemisphere damage, the left-to-right mode of processing that emanates from the right hemisphere is presumably disrupted, along with other consequences that may affect all spatial processing to the left of midline. When the right hemisphere is damaged, the reading errors that preferentially degrade left-most letter positions might reasonably be regarded as a form of left neglect. In contrast, the conclusion we are compelled to reach is that the pattern of error at issue here – word substitutions with degraded right-most letter positions – should not be regarded as a type of neglect but as a form of language impairment.

Acknowledgements

The studies reported here were supported by grant number R01-DC00699 from the National Institute on Deafness and Other Communication Disorders to the University of Maryland School of Medicine. The authors are grateful to James A. Reggia for helpful comments, and to Anne N. Haendiges for assistance with the patient data.

References

Berndt, R.S., Haendiges, A.N., Mitchum, C.C., Wayland, S.C. (1996) An investigation of nonlexical reading impairments. Cog. Neuropsychol., 13: 763-801.

Berndt, R.S., Haendiges, A.N. (1997) Positional effects in dyslexic "visual errors": constraints on the interpretation of word substitutions. Brain Lang., 60: 112-115.

Berry, M.J., Warland, D.K., Meister, M. (1997) The structure and precision of retinal spike trains. Proc. Natl. Acad. Sci. USA, 94: 5411-5416.

Buxbaum, L.J., Coslett, H.B. (1996) Deep dyslexic phenomena in a letter-by-letter reader. Brain Lang., 54: 136-167.

Caramazza, A., Hillis, A.E. (1990a) Levels of representation, co-ordinate frames, and unilateral neglect. Cognitive Neuropsychology, 7: 391-445.

Caramazza, A., Hillis, A.E. (1990b) Spatial representation of words in the brain implied by the studies of a unilateral neglect patient. Nature, 346: 267-269.

Casey, T., Ettlinger, G. (1960) The occasional "independence" of dyslexia and dysgraphia from dysphasia. J. Neurol. Neurosurg. Psychiat., 23: 228-236.

Chambers, S.M. (1979) Letter and order information in lexical access. Journal of Verbal Learning and Verbal Behavior, 18: 225-241.

Chokron, S., De Agostini, M. (1995) Influence of reading habits on line bisection: a developmental approach. Cog. Brain Res., 3: 51-58.

Chokron, S., Imbert M. (1993) Reading Habits and line bisection. Cog. Brain Res., 1: 219-222.

Coltheart, M. (1980) Deep dyslexia: a review of the syndrome. In M. Coltheart, K. Patterson and J.C. Marshall (Eds.), Deep Dyslexia. Routledge and Kegan Paul, London, pp. 22-47.

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

De Ruyter van Steveninck, R.R., Lewen, G.D., Strong, S.P., Koberle, R., Bialek, W. (1997) Reproducibility and variability in neural spike trains. Science, 275: 1805-1808.

Ellis, A.W., Young, A.W., Flude, B.M. (1993) Neglect and visual language. In I.H. Robertson and J.C. Marshall (Eds.), Unilateral Neglect: Clinical and Experimental Studies. Lawrence Erlbaum, Hillsdale, pp. 233-256.

Eng, T.L., Hellige, J.B. (1994) Hemispheric assymetry for processing unpronounceable and pronounceable letter trigrams. Brain Lang., 46: 517-535.

Ferrand, L., Grainger, J. (1993) The time course of orthographic and phonological code activation in the early phases of visual word recognition. Bull. Psychonomic Soc., 31: 119-122.

Grainger, J., Jacobs, A.M. (1991) Masked constituent letter priming in an alphabetic decision task. Eur. J. Cog. Psychol., 3:413-434.

Greenwald, M.L., Berndt, R.S. (1998) Impaired encoding of abstract letter order: severe alexia in a mildly aphasic patient. Cognitive Neuropsychology, in press.

Hammond, E.J., Green, D.W. (1982) Detecting targets in letter and non-letter arrays. Canadian Journal of Psychology, 36: 67-82.

Heilman, K.M. (1995) Attentional asymmetries. In R.J. Davidson and K. Hugdahl (Eds.), Brain Asymmetry. MIT Press, Cambridge, pp. 217-234.

Hellige, J.B. (1993) Hemispheric Asymmetry: What's Right and What's Left. Harvard University Press, Cambridge.

Hellige, J.B., Cowin, E.L., Eng, T.L. (1995) Recognition of CVC syllables from LVF, RVF, and central locations: hemispheric differences and interhemispheric interaction. J. Cog. Neur., 7: 258-266.

Hillis, A.E., Caramazza, A. (1995) A framework for interpreting distinct patterns of hemispatial neglect. Neurocase, 1: 189-207.

Hopfield, J.J. (1995) Pattern recognition computation using action potent ial timing for stimulus representation. Nature, 376: 33-36.

Hopfield, J.J. (1996) Transforming neural computations and representing time. Proc. Natl. Acad. Sci. USA, 93: 15440-15444.

Humphreys, G.W., Evett, L.J., Quinlan, P.T. (1990) Orthographic processing in visual word identification. Cog. Psychol., 22: 517-560.

Joliot, M., Ribary, U., Llinas R. (1994) Human oscillatory brain activity near 40 Hz coexists with cognitive temporal binding. Proc. Natl. Acad. Sci. USA, 91:11748-11751.

Kinsbourne, M. Warrington, E.K. (1962) A variety of reading disability associated with right hemisphere lesions. J. Neurol. Neurosurg. Psychiat. 25: 339-344.

Lefton, L.A., Fisher, D.F., Kuhn, D.M. (1978) Left-to-right processing of alphabetic material is independent of retinal location. Bulletin of the Psychonomic Society, 112: 171-174.

Lisman, J.E., Idiart, M.A.P. (1995) Storage of 7 +- 2 short-term memories in oscillatory subcycles. Science, 267: 1512-1515.

Llinas, R., Ribary, U. (1993) Coherent 40-hz oscillation characterizes dream state in humans. Proc. Natl. Acad. Sci. USA, 90: 2078-2081.

Mason, M. (1982) Recognition time for letters and nonletters: effects of serial position, array size, and processing order. J. Exp. Psychol., 8: 724-738.

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

Mewhort, D.J.K., Beal, A.L. (1977) Mechanisms of word identification. J. Exp. Psychol. 3: 629-640.

Montant, M., Nazir, T.A., Poncet, M. (1998) Pure alexia and the viewing position effect in printed words. Cog. Neuropsychol., 15: 93:140.

Morton, J., Patterson, K. (1980) A new attempt at an interpretation, or, an attempt at a new interpretation. In M. Coltheart, K. Patterson and J.C. Marshall (Eds.), Deep Dyslexia, Routledge & Kegan Paul, London, pp. 91-118.

Mozer, M.C. (1987) Early parallel processing in reading: a connectionist approach. In M. Coltheart (Ed.), Attention and Performance XII: The Psychology of Reading. Lawrence Erlbaum, London, pp. 83-104.

Mozer, M.C., Behrmann, M. (1992) On the interaction of selective attention and lexical knowledge: a connectionist account of neglect dyslexia. J. Cog. Neurosci. 2: 96-123.

Perea, M. (1998) Orthographic neighbors are not all equal: evidence using an identification technique. Language and Cognitive Processes 13: 77-90.

Peressotti, F., Grainger, J. (1995) Letter-position coding in random consonant arrays. Perception & Psychophysics, 57: 875-890.

Peressotti, F., Grainger, J. The Role of Letter Identity and Letter Position in Orthographic Priming. Perception & Psychophysics, in press.

Perfetti, C.A., Bell, L. (1991) Phonemic activation during the first 40 ms of word identification: evidence from backward masking and priming. J. Mem. Lang., 30: 473-485.

Plaut, D.C., Shallice, T. (1993) Deep dyslexia: a case study of connectionist neuropsychology. Cog. Neuropsychol., 10:377-500.

Proctor, R.W., Healy, A.F. (1985) Order-relevant and order-irrelevant decision rules in multiletter matching. J. Exp. Psychol., 11: 519-537.

Proctor, R.W., Healy, A.F. (1987) Task-specific serial position effects in comparisons of multiletter strings. Perception & Psychophysics, 42: 180-194.

Ratcliff, R. (1981) A theory of order relations in perceptual matching. Psychol. Rev., 88: 552-572.

Reuter-Lorenz, P.A., Baynes, K. (1992) Modes of lexical access in the callosotomized brain. J. Cog. Neur., 4: 155-164.

Reuter-Lorenz, P.A., Brunn, J.L. (1990) A prelexical basis for letter-by-letter reading: a case study. Cog. Neuropsychol., 7: 1-20.

Riddoch, J. (1990) Neglect and the peripheral dyslexias. Cog. Neuropsychol. 7: 369-389.

Rieke, F., Warland, D., De Ruyter van Steveninck, R., Bialek, W. (1997) Spikes: Exploring the Neural Code, MIT Press, Cambridge.

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

Sejnowski, T.J., Rosenberg, C.R. (1987) Parallel networks that learn to pronounce English text. Complex Systems, 1: 145-168.

Shallice,T., Warrington, E.K. (1975) Word recognition in a phonemic dyslexic patient. Quarterly Journal of Experimental Psychology, 27: 187-199.

Shallice, T. (1988) From Neuropsychology to Mental Structure. Cambridge University Press, New York.

Tiitinen, H., Sinkkonen, J., Rainikainen, K., Alho, K., Lavi-kainen, J., Naatanen, R. (1993) Selective attention enhances the 40-hz response in humans. Nature, 364: 59-60.

Victor, J.D., Purpura, K.P. (1996) Nature and precision of temporal coding in visual cortex: a metric-space analysis. J. Neurophysiol., 76: 1310-1326.

Warrington, E.K., Zangwill, O.L. (1957) A study of dyslexia. J. Neurol. Neurosurg. Psychiat. 20: 208-215.

Warrington, E.K. (1991) Right neglect dyslexia: a single case study. Cog. Neuropsychol.,, 8: 193-212.

Whitely, H.E., Walker, P. (1994) The activation of mulitletter units in visual word recgnition, Visual Cognition, 1: 433-473.

Whitely, H.E., Walker, P. (1997) Mulitletter units in visual word recgnition: direct activation by supraletter features, Visual Cognition, 4: 69-110.

Whitney, C.S., Berndt, R.S., Reggia, J.A. (1996) Simulation of neurogenic reading disorders with a dual-route connectionist model. In J.A. Reggia, E. Ruppin, R.S. Berndt (Eds.), Neural Modeling of Brain and Cognitive Disorders, World Scientific, Singapore, pp. 201-228.

Appendix A – Temporal Encoding Simulations

Entity	Definition
[illegible]	[illegible]

The functions and variables defined in the model are summarized above. To review, prior to firing, $V_i(t) = E_i + C(t) + I_i(t) + v$, where t ranges from 0 to g. When $V_i(t) \geq T$, L_i fires and sends lateral inhibition to all other letter nodes. After firing, V_i falls to v, and then rebounds, where the rate of rebound depends on E_i.

The simulations were performed using the following values: $T = -50$, $g = 200$, and $v = -65$. The values E_i constituted the inputs to the simulations. For simplicity, C was taken to be piecewise linear, with $C(t) = 0.1 * t$ for $t \leq g/2$, and $C(t) = C(g/2) - C(t - g/2)$ for $t > g/2$. If a node has fired, $V(t) = V(t-1) + T/(T - E_i + 1)$ up to a maximal value of $v + C(t) + E_i$. For simplicity, the effect of lateral inhibition on a node was considered to be negligible, unless the node had already fired, in which case the node was inhibited for the rest of the cycle.

Appendix B – String Comparison Simulations

Using a trial and error process, we arrived at parameters that yielded a good fit to the experimental data. A waveform encodes position within a string, and is comprised of time slots 1 to 120. The representation of a string consists of four letters and their positional waveforms. A value is associated with each time slot in a waveform. For simplicity, values are abstract, and limited to three values: 10 while firing, -1 prior to firing, and 4 after firing. The first letter in a string fires during the first 50 slots, the second letter fires during the next 25 slots, the third letter during the next 13 slots, and the fourth and final letter fires during the next 25 slots. We denote the waveform representing position i as W_i, and we denote the value of the kth time slot of W_i as $P(W_i, k)$. For example, for the waveform representing the second position, $P(W_2, 1)$ to $P(W_2, 50) = -1$, $P(W_2, 51)$ to $P(W_2, 75) = 10$, and $P(W_2, 76)$ to $P(W_2, 120) = 4$. If a letter does not fire, its positional waveform, W_0, has value -1 for all time slots.

We assume that a mechanism exists for comparing firing patterns (waveforms). Briefly, the similarity of two strings is evaluated by comparing the waveforms for each letter, calculating a total raw score, and comparing the raw score to the maximal possible score. More specifically, we define a comparsion function, $C(W_i, W_j)$, which takes two positional waveforms and yields a comparison waveform, W_{ij}. C is defined as follows: For each time slot t, if $P(W_i, t) * P(W_j, t) > 0$, $P(W_{ij}, t) = (P(W_i, t) + P(W_j, t))/2$; otherwise, $P(W_{ij}, t) = P(W_j, t) - 6$. We define $L(S_1, S_2, i)$ as the position in S_2 of the ith letter of S_1; if the letter does not appear in S_2, $L = 0$. To compare two strings, the comparison waveforms for each letter are calculated. That is, given strings S_1 and S_2, for each position i, we calculate $C(W_i, W_{L(S_1, S_2, i)})$. For example, to compare the strings $S_1 = \text{ABCD}$ and $S_2 = \text{AXBC}$, the following positional waveforms are compared: W_1 and W_1 (for A), W_2 and W_3 (for B), W_3 and W_4 (for C), and W_4 and W_0 (for D). The raw score, $R(S_1, S_2)$, is calculated by summing over all time slots over all four comparison waveforms. The maximal raw score, M, is set to the $R(S_1, S_1)$. The simulated response time, $T(S_1, S_2)$ is calculated by comparing $R(S_1, S_2)$ to M and scaling it; $T(S_1, S_2) = 800\text{ms} * (R(S_1, S_2)/M)^3 + 720\text{ms}$.

Appendix C – Activation Gradient Simulations (II)

To represent the input string, activations of bigrams comprised of contiguous letters in the input string were set to 0.6^{pos-1}, where *pos* denotes the position of the initial letter of the bigram. For positions greater than 3, *pos* was set to 3. Thus, these activation levels were either 1.0, 0.6 or 0.36. The activations of non-contiguous bigrams were set to 0.6^{pos}. If a bigram occurred more than once, its activation was set to the sum of the activations for each occurrence.

We assumed that despite the separation between the initial letter and the final letter, their high levels of activation would combine to make bigram B_{1n} highly activated. We also assumed that the activation of the final letter depends on the length of the input string, because for longer input strings, the final letter starts to fire later in the cycle and thus achieves a lower level of activation. To approximate these assumptions, the activation of B_{1n} was adjusted to $1.0 - 0.01n$. For example, to represent the input string FLANK, FL = 1.0, LA = 0.6, AN = 0.36, NK = 0.36, FA = 0.6, LN = 0.36, AK = 0.22, FN = 0.6, LK = 0.36, FK = 0.95, and all other bigrams were set to 0. Two-letter and one-letter words were handled as special cases.

Figure Captions

Figure 1: Architecture of the model. At the letter level, simultaneous graded inputs and lateral inhibition create a temporal firing pattern, as indicated by the timing of firing displayed under the letter nodes. Excitatory connections link the letter nodes and the bigram nodes, which recognize ordered pairs of letters. The activations of the bigrams (shown above the nodes) are proportional to the activations of the constituent letters. Excitatory connections link the bigram nodes and the word nodes. Activation of word nodes is based on the conventional dot-product model.

Figure 2: Simulated firing at the letter level. Each point indicates a time at which the corresponding node spiked. Node 1 received the highest amount of input, while node 5 received the lowest. The upper graph shows the results for a high overall level of input, E = [15 13.5 12.0 10.5 9.0]. The first node fires early in the cycle ($t = 0$), and the last node fires more than the internal nodes (5 spikes versus 2 or 3 spikes). For a lower level of input, E = [13.5 11.5 9.5 7.5 5.5], displayed in the lower graph, the pattern is different. The first node fires later in the cycle ($t = 15$), and the last node fires only once, having the lowest level of activation.

Figure 3: Experimental and simulated results for the order task of the string comparison study. The upper graph dislays the replacement pairs by position of replacement. The lower graph displays the permutation pairs. Permutations are ordered by degree of positional displacement of the constituent letters from the base string, ABCD.

Figure 4: Retention level by position for experimental and simulated results.

Figure 5: Retention level for position three by target length for experimental and simulated results.

Figure 6: Average error length by target length for experimental and simulated results.

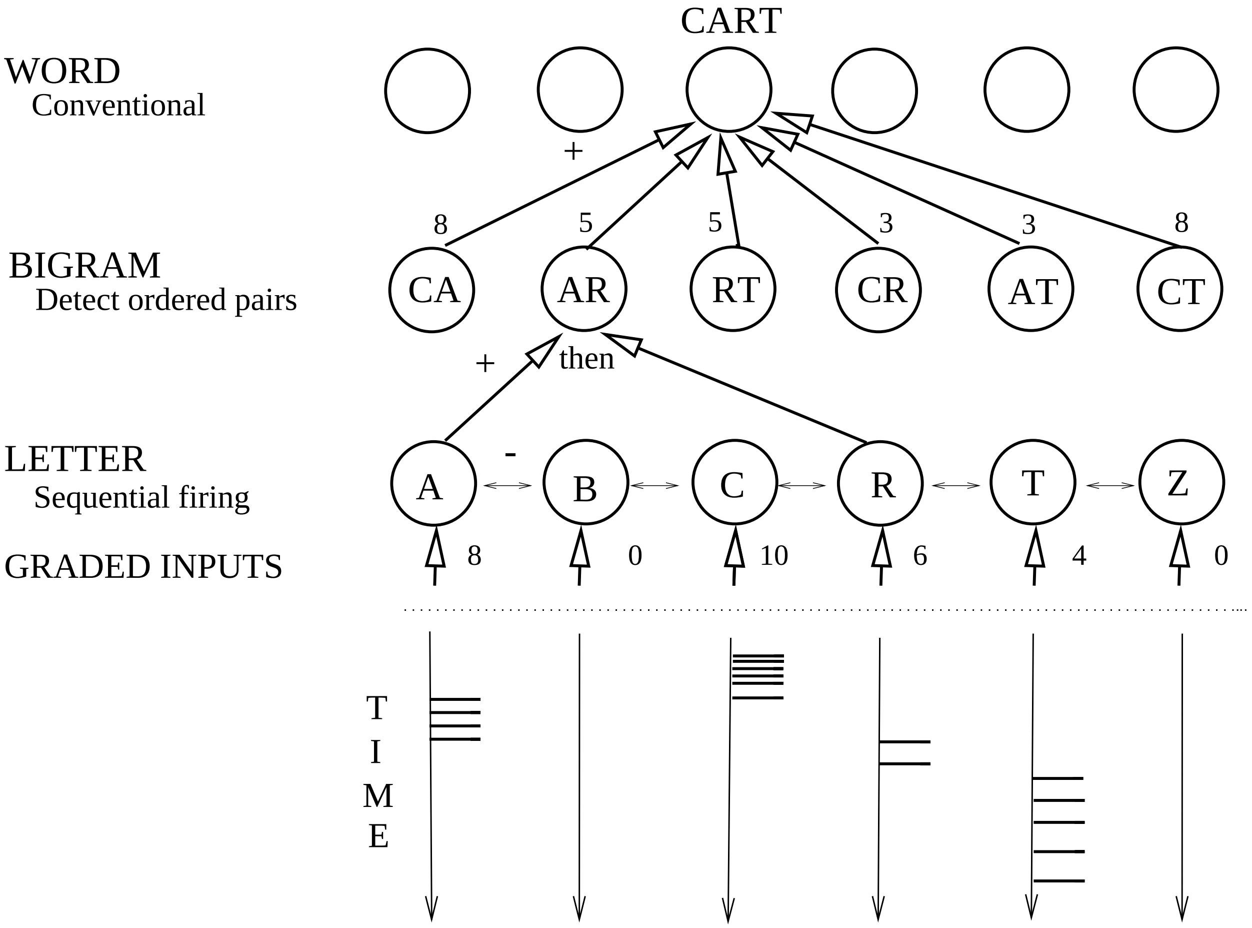
CART
WORD
Conventional
+
8
5
5
3
3
8
BIGRAM
Detect ordered pairs
CA
AR
RT
CR
AT
CT
+
then
LETTER
Sequential firing
-
A
B
C
R
T
Z
GRADED INPUTS
8
0
10
6
4
0
T
I
M
E

FIGURE 1

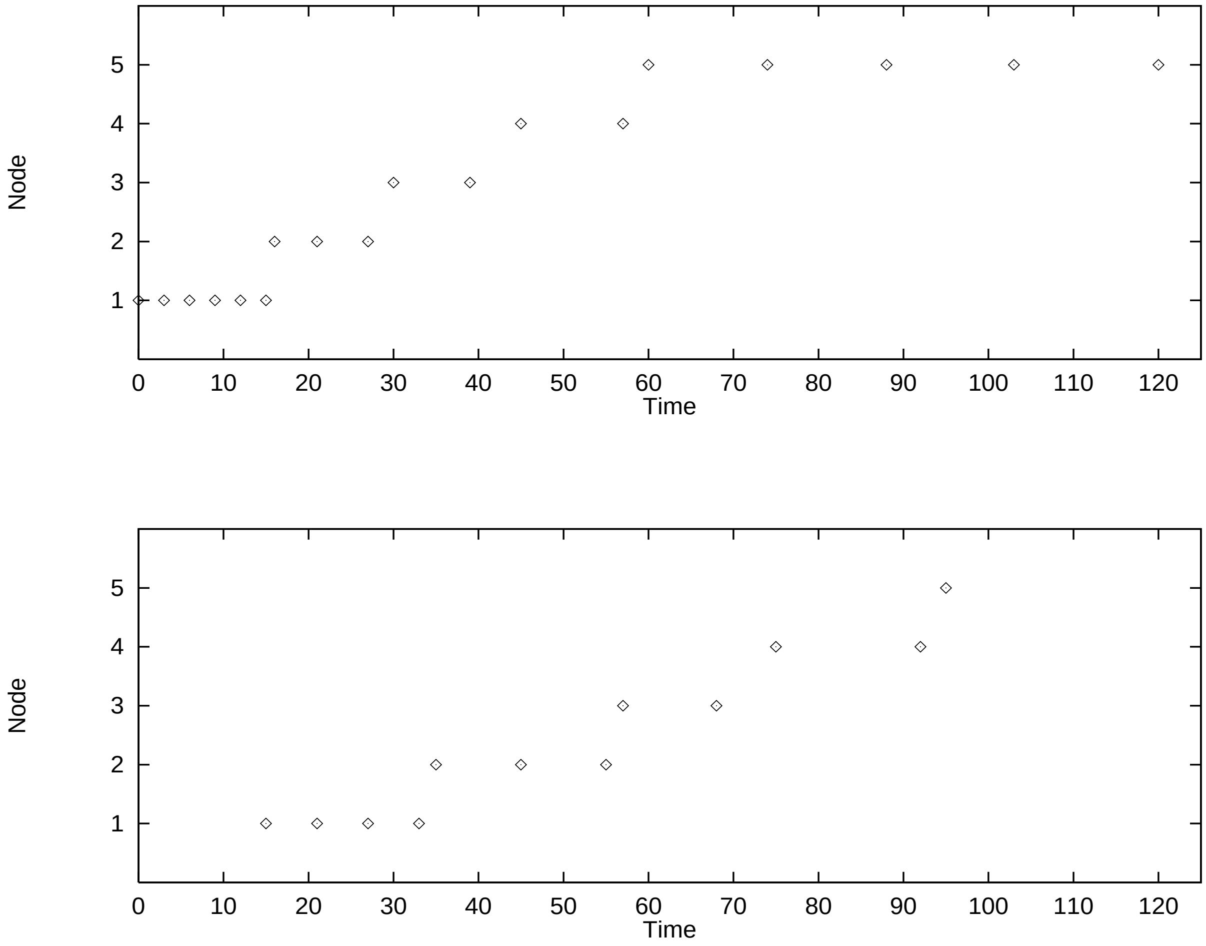

FIGURE 2

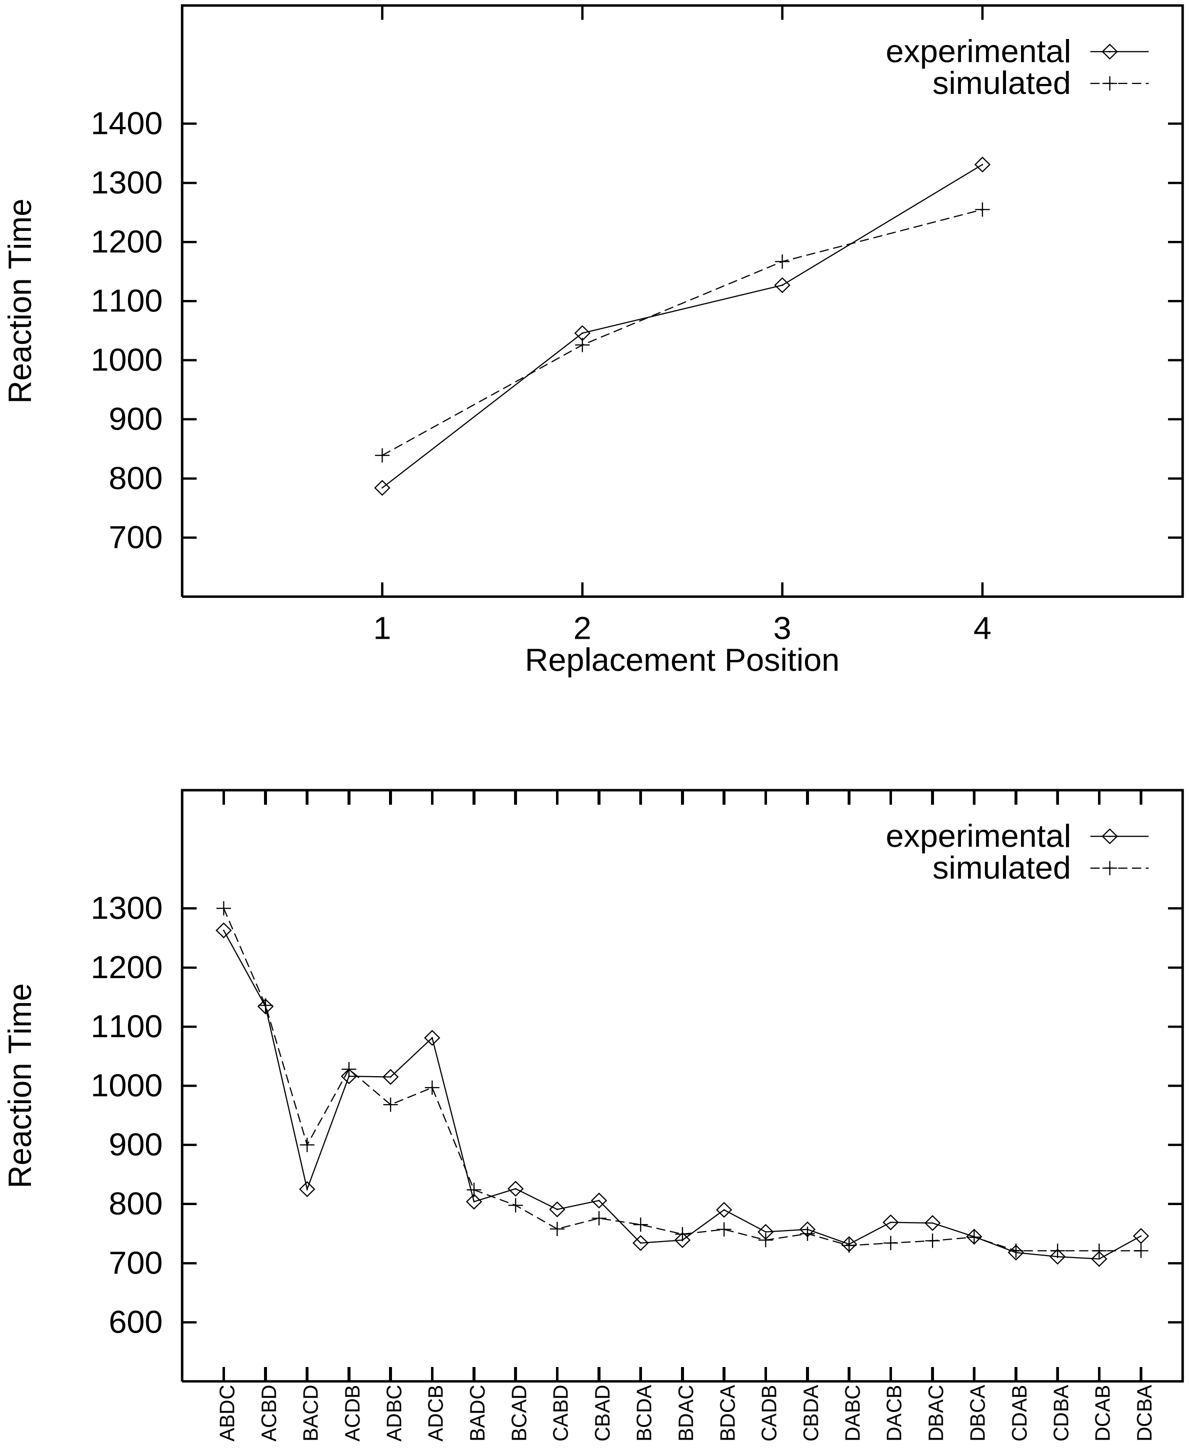
experimental
simulated
Reaction Time
1400
1300
1200
1100
1000
900
800
700
1
2
3
4
Replacement Position
experimental
simulated
Reaction Time
1300
1200
1100
1000
900
800
700
600
ABDC
ACBD
BACD
ACDB
ADBC
ADCB
BADC
BCAD
CABD
CBAD
BCDA
BDAC
BDCA
CADB
CBDA
DABC
DACB
DBAC
DBCA
CDAB
CDBA
DCAB
DCBA

FIGURE 3

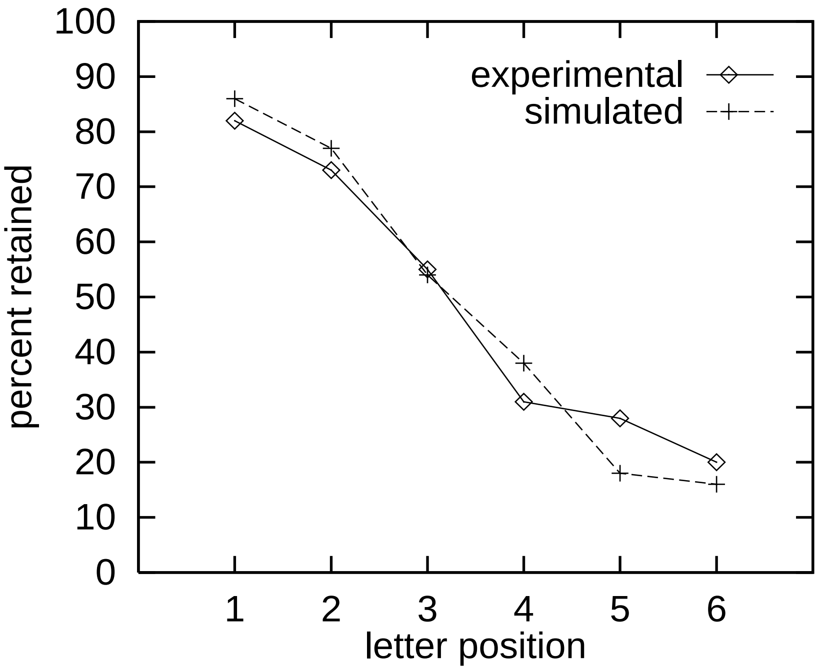

100
90
80
70
60
50
40
30
20
10
0
percent retained
1
2
3
4
5
6
letter position
experimental
simulated

FIGURE 4

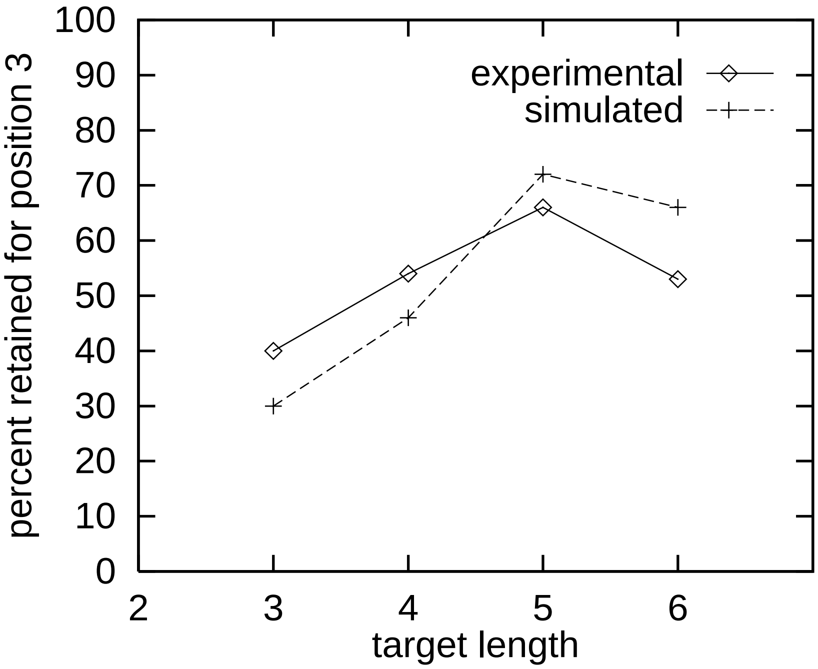
percent retained for position 3
100
90
80
70
60
50
40
30
20
10
0
2
3
4
5
6
target length
experimental
simulated

FIGURE 5

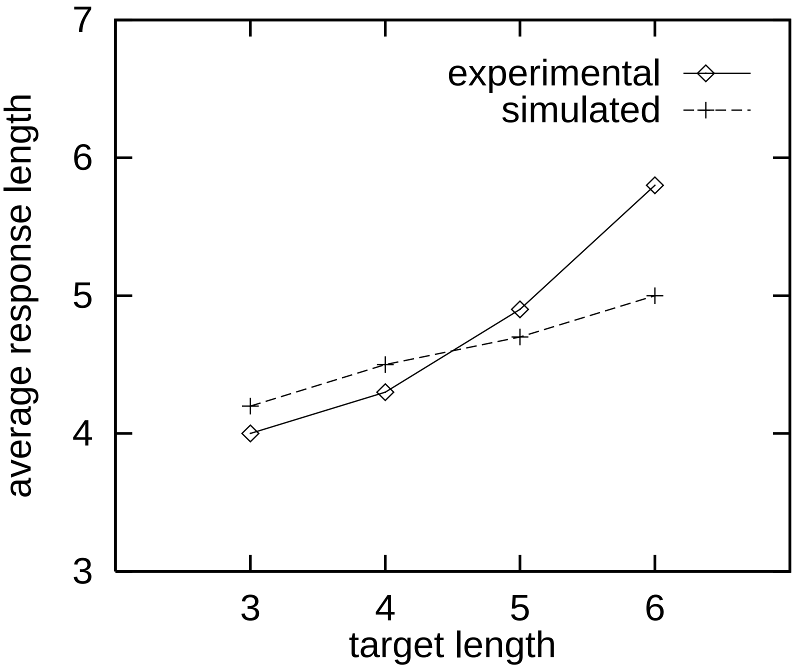
7
6
5
4
3
average response length
experimental
simulated
3
4
5
6
target length

FIGURE 6