Complex Systems 2 (1988) 509–547

Competitive Dynamics in a Dual-route Connectionist Model of Print-to-sound Transformation

James A. Reggia
Department of Neurology, University of Maryland Hospital, 22 South Greene Street, Baltimore, MD 21201, USA
and
Department of Computer Science, University of Maryland Institute for Advanced Computer Studies, College Park, MD 20742, USA

Patricia M. Marsland
Rita Sloan Berndt
Department of Neurology, University of Maryland Hospital, 22 South Greene Street, Baltimore, MD 21201, USA

Abstract. This paper describes a connectionist model of print-to-sound transformation ("word naming" or "reading aloud"). The associative network it uses is based on published studies of oral reading, and simulation results are compared to experimental data in the psychological literature. The results obtained are of interest for two separate reasons. First, the print-to-sound connectionist model is based on an *indirectly interactive dual-route hypothesis* of reading aloud. The model confirms that this hypothesis, when implemented as a detailed and sizeable computer simulation, can account qualitatively for a number of behavioral phenomena such as regularity and word frequency effects. The model thus provides support for a modified dual-route hypothesis involving indirectly interactive routes and verifies the hypothesis' consistency with a set of replicable psychological data. The second reason the print-to-sound connectionist model is of interest is that it uses a new approach to implementing competitive dynamics in connectionist models. Focused spread of network activation and avoidance of network saturation are produced by using a *competitive activation mechanism* rather than explicit inhibitory links between competing nodes. The print-to-sound model demonstrates for the first time that competitive activation mechanisms can function usefully in relatively large, complex situations of interest in cognitive psychology and artificial intelligence.

1. Introduction

The nature of the cognitive mechanisms underlying pronunciation of printed text is currently a dynamic and controversial topic. One of the most influential theories about these mechanisms, the dual-route hypothesis, proposes that reading aloud can be achieved by either of two distinct procedures: a *lexical* procedure by which letter strings are associated with whole-word pronunciations, and a *non-lexical* procedure by which segments of printed words (i.e., letters or letter clusters) are associated with sub-lexical sound patterns, and then "assembled" into a pronunciation. The standard version of this hypothesis, referred to as the *independent dual-route hypothesis* in this paper, starts with the assumption that these two information processing "routes" are independent (e.g., [8,35]). While this independence assumption is perhaps the simplest perspective one can take, a convincing case can be made that the *independent* dual-route hypothesis has difficulty accounting for existing empirical data (e.g., [30]). Many of the criticisms that have been directed at the dual-route model can be countered, and the basic architecture of the model retained, if the assumption is relinquished that the two information processing routes involved are independent [15]. For this solution to be tenable, however, it becomes necessary to define clearly and explicitly the nature of the interactions between the lexical and non-lexical routes, and to provide some support for the viability of such an approach.

These considerations motivated the development of a connectionist model of print-to-sound transformation ("word naming," "reading aloud") based on an *indirectly interactive dual-route hypothesis*. This print-to-sound connectionist model uses a dual-route associative network to represent correspondences between relevant linguistic units. The term "indirectly interactive" is used to characterize the model because, although the flow of activation through the lexical and non-lexical routes is not completely independent, these two routes influence one another only in a limited, indirect fashion.

The behavior of the computational model described in this paper has been examined systematically for word frequency and regularity effects and has been shown to exhibit behavior qualitatively similar to that seen with normal readers. Limited studies with "lesioned" versions of the model also demonstrate behavior having similarities to empirical observations of readers with acquired dyslexia. The print-to-sound connectionist model thus provides support for an indirectly interactive dual-route hypothesis by demonstrating that, when made explicit in a detailed computer simulation, this hypothesis is consistent with some important results obtained in empirical studies of normal and dyslexic readers.

The print-to-sound model is also of interest in that it uses a new method for controlling spreading activation in connectionist models. This method is referred to as a *competitive activation mechanism*. In fact, it is this approach to controlling network dynamics that forms the basis for the limited, indirect interactions between the two routes in the print-to-sound model. To understand the issues raised by competitive activation mechanisms, it is important to recognize that connectionist models like the print-to-sound system (i.e., those using an associative network to represent memory) are generally not "neural models," even though they use a neuron-like network and processing paradigm. Their nodes and connections represent concepts and associations rather than neurons and synapses. They typically model functional cognitive mechanisms rather than biophysical brain processes, and are studied primarily by workers in cognitive science and artificial intelligence (AI) rather than neuroscience [44].

In spite of these distinctions, many associative network models have adopted methods for controlling spreading activation that were initially intended for modelling networks of biological neurons. For example, competitive interactions between nodes intended to produce a winner-takes-all behavior have usually been implemented through the use of lateral inhibitory links [14,20,34,57]. As explained below (section 3), such an approach to focusing spread of activation involves significant problems when applied in associative rather than neural networks.

Recently, a different approach to introducing competitive interactions into connectionist models of associative memory has been proposed [40,41]. Rather than implementing direct competitive behavior through explicit *structural* features of a network (inhibitory links), competition is introduced into the *functional* mechanism or rule by which the spread of activation is controlled. The print-to-sound model demonstrates for the first time that such competitive activation mechanisms can function effectively in connectionist models having networks of a size and complexity often found in contemporary cognitive science and AI systems.

The dual-route print-to-sound transformation was selected as the first large-scale test of competitive activation mechanisms both because it provides an implementation of indirectly interactive routes, and because print-to-sound transformation is a challenging but relatively well-defined and circumscribed problem. The print-to-sound transformation is challenging in that the existence of two hypothesized "routes" by which information flows through the underlying associative network implies that phoneme nodes serving as outputs for the model might receive conflicting information about what their activation levels should be. Such conflicts must be resolved by these phoneme nodes based on locally available information as node activations approach equilibrium. On the other hand, compared to many other cognitive tasks involving associative memory (e.g., natural language processing at the semantic level), the underlying network is relatively circumscribed; hence, a fairly significant subset of it can be captured in a model. There is also a relatively large amount of empirically-derived information in the literature upon which to base network structure, assignment of weights to connections, and analysis of simulation results.

The remainder of this paper is organized as follows. Section 2 summarizes empirical studies of print-to-sound transformation and the dual-route hypothesis. The associative network used in the connectionist model to represent the relevant correspondences between graphemes, phonemes, and words is described. Section 3 discusses competitive behavior in connectionist models and explains some difficulties in implementing competition with lateral inhibitory connections when a local representation of information is used (as in the print-to-sound model). The specific competitive activation mechanism used in the print-to-sound model is summarized, and the notions of independent and interactive routes are explained. Section 4 describes a limited but systematic study of the print-to-sound model's behavior. Section 5 concludes by summarizing this work and its implications.

2. Print-to-sound transformation

The term *print-to-sound transformation* is used here to refer to the task of reading aloud a single printed word. The specifics of the connectionist model of print-to-sound transformation described in this paper are strongly influenced by the *dual-route hypothesis* of reading that postulates the existence of two parallel and independent "routes" of information flow during reading [1,8,9,35]. We first briefly review this hypothesis and then describe the associative network used in the print-to-sound connectionist model.

2.1 Empirical studies of the print-to-sound mapping

During the last several years there has been a great deal of empirical research on the cognitive processes underlying skilled reading [13,19]. A significant part of this research has focused on oral reading or "word naming." In performing the task of reading a word aloud, one transforms a sequence of printed graphemes into a sequence of spoken phonemes. A *grapheme* is defined to be one to a few printed characters serving as the written representation of a phoneme (following [10]). For typographic convenience, phonemes will be represented by lower-case letters between slanted lines, graphemes by upper-case letters, and word/morphemes by double-quote marks. For example, the word "onion" consists of five graphemes, O N I O N, which correspond to five phonemes, /uh$^+$ n y uh- n/.[1]

The cognitive task of word naming or reading aloud is often represented diagramatically as information flow and transformation through a number of cognitive modules involving two routes (e.g., [19,36,48]). Figure 1 provides a simplified example of such a diagram involving two routes by which information flows as it is transformed from written to spoken form. One route is the *grapheme-phoneme correspondence or GPC route* (bottom of figure 1). Reading aloud via this process involves mapping graphemes onto their corresponding phonemes. A major reason for postulating such a route is the ability of skilled readers to read aloud pronounceable non-words like "kint." Although other possible mechanisms for non-word reading have been postulated (e.g., [32,23]), these lack sufficient specificity for adequate testing.

[1]Symbols used here are keyboard-compatible representations of the International Phonetic Alphabet. The phonemes they represent are largely obvious in this paper, but a complete listing and definitions can be found in [4].

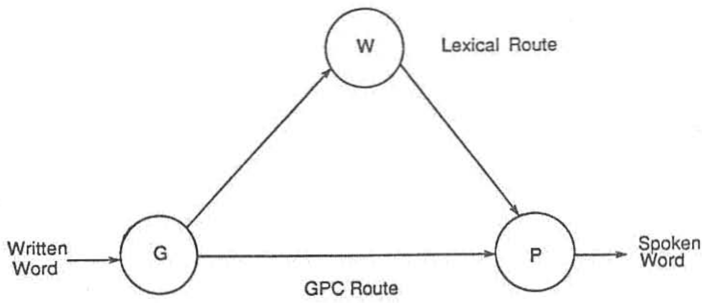

Figure 1: A simplified characterization of the Print-to-Sound Mapping. G = graphemes, P = phonemes, W = words. The GPC route is along the bottom (G $\rightarrow$ P), the lexical route along the top (G $\rightarrow$ W $\rightarrow$ P). See text for details.

A second route used in reading aloud will be referred to as the *lexical route* (top of figure 1). Reading aloud via this route involves the mapping of graphemes onto morphemes/words via the visual word recognition system and then onto their phonemic representation. This route is postulated to account for the ability of skilled readers to correctly pronounce exception ("irregular") words and to distinguish the meaning of distinct words that happen to sound the same (e.g., "weak" and "week"; see [27, ch. 5]).

The dual-route model of reading aloud is strongly supported by additional evidence from studies of individuals with acquired dyslexia (alexia). Acquired dyslexia is a disorder of reading resulting from brain damage in formerly literate readers. From a linguistic point of view, a variety of forms of acquired dyslexia exist, and form-specific features can be correlated with localized functional impairments in information flow models like that of figure 1. For example, in phonological dyslexia patients can read many familiar words but usually fail to read aloud even simple non-words. In such cases it has been postulated that a selective impairment of the GPC route has occurred [3,22]. In contrast, surface dyslexics can read non-words and real words with "regular" spelling patterns, but fail to read most "exception" words (e.g., "yacht") where grapheme-phoneme correspondences do not yield the correct pronunciation [13]. Their responses to such irregular words are typically "regularizations"; e.g., reading "yacht" as /y ae ch t/. These patients appear to suffer from impairment to the lexical route, with relative sparing of the GPC route. The important fact that both of these syndromes have been described (a "double dissociation" of symptoms) indicates that the two routes are functionally separate, and are not simply extremes on a continuum of processing resource with one type of process more susceptible to the effects of brain damage than the other. Such relatively selective functional deficits have generated a great deal of research and interest among cognitive psychologists during the last few years [12,19,36,48].

Both the lexical and the GPC routes have been implicated in the normal pronunciation of words by skilled readers. In research measuring the time taken for normal subjects to read words aloud (pronunciation latency), it has been found, in general, that words with regular pronunciations can be read faster than irregular words [24,53]. This finding has been interpreted as an indication that congruency in the output of the two routes yields faster response times. A complication was subsequently added to this view by the finding that this effect of regularity on pronunciation latency was detectable only for words of low frequency [49]. High-frequency words, whether regular or irregular, were pronounced uniformly quickly relative to low-frequency words. This finding suggests that any simple "horse-race" model, in which the two parallel and independent routes are used to achieve the same end under different conditions, underestimates the complexity of the reading process. This result, which appears to require some kind of flexible but structured interaction between the two routes, is the primary focus of the print-to-sound model discussed here.

2.2 Associative network structure

The network in the connectionist model of print-to-sound transformation uses a local representation of information. Nodes represent graphemes, phonemes and words, while connections represent positively weighted associations between these entities. The overall network structure is illustrated in Figure 1, where each pictured oval represents a set of node types, and each pictured arc represents numerous forward connections. It can be seen that there are two *routes* by which activation can flow through the network: the "lower" GPC route and the "upper" lexical route. Running a simulation involves selective application of an externally supplied source of input to the appropriate grapheme nodes, thereby driving up their activation levels. Activation then spreads from grapheme nodes to phoneme nodes (via the GPC route), and from grapheme nodes to "hidden" word nodes to phoneme nodes (via the lexical route). The activation levels of phoneme nodes represent the network's output.

Since each word node in such a network connects to multiple grapheme/phoneme nodes occurring in specific positions, there are actually multiple copies or *instances* of grapheme and phoneme node sets in any simulation. The exact number of instances of grapheme and phoneme node sets is determined by the number of graphemes that are designated as input. For example, if a specific simulation involved the presentation of a sequence of n graphemes as input, this is implemented in the model by dynamically constructing n copies of the grapheme nodes prior to initiating the simulation. Each set of grapheme nodes corresponds to one input position, where positions are numbered from 1 (initial position) to n (final position). For each set of grapheme nodes so constructed, a corresponding set of phoneme nodes is generated (recall that a grapheme is defined as the orthographic representation of a single phoneme) along with all relevant connections from graphemes to phonemes in one position. Thus, for an input sequence of n graphemes, counting the set of word nodes there are $2n + 1$ interconnected sets of nodes present. This use of duplicate but position-specific sets of nodes is similar to that used in other connectionist models of "low-level" linguistic information processing [18,34].

As an example, suppose the connectionist model is presented with the five graphemes

O N I O N

as input. The network constructed for the simulation would contain five sets of grapheme nodes and five sets of phoneme nodes as illustrated in figure 2. In what follows, all of the nodes and connections occurring in a single grapheme/phoneme position are referred to as forming a *path* through the network. Thus, in figure 2 all of the G_3 and P_3 nodes and their connections are the third of five paths forming the GPC route. Similarly, the G_3, W and P_3 nodes and their connections are one of five paths forming the lexical route. The individual paths in the GPC route are separate from one another in this model, while those in the lexical route are not since they converge at the single set of word nodes.

There are 48 phonemes and 168 graphemes represented as nodes in *each* position-specific path in the print-to-sound model's associative net work. These nodes and their connections, which form the GPC route, are based on data from the analysis of a corpus of 17,310 words [26]. That study defined graphemes as letters or letter clusters corresponding to a single phoneme, using a one-to-one correspondence between graphemes and phonemes in words. Motivated by educational issues related to spelling, the study by Hanna et al. provided for any given phoneme a list of its possible spellings (graphemes) and their frequencies [26, tables 17 and 18]. The information needed in the print-to-sound model network is the "reverse" of this available sound-to-print information, which cannot be directly retrieved from individual table entries in the source document. For example, although the phoneme /aw/ is only occasionally written as AU (probability = .15), the grapheme AU is almost always pronounced as /aw/ (probability = .95).

For this reason, a computer program was implemented to generate the grapheme-to-phoneme connections and weights (conditional probabilities) needed to form the GPC route in the print-to-sound connectionist model. This program used a slightly revised version of the tables in the source document [26]. A listing of the resulting grapheme-to-phoneme associations and their weights as used in our model, as well as the details of their derivation, can be found in [4]. An example of a single grapheme node and its connections to phoneme nodes in a path of the GPC route is illustrated in figure 3. For the special cases of graphemes in the first and final position of a word, modified weights were derived in the same fashion [26].

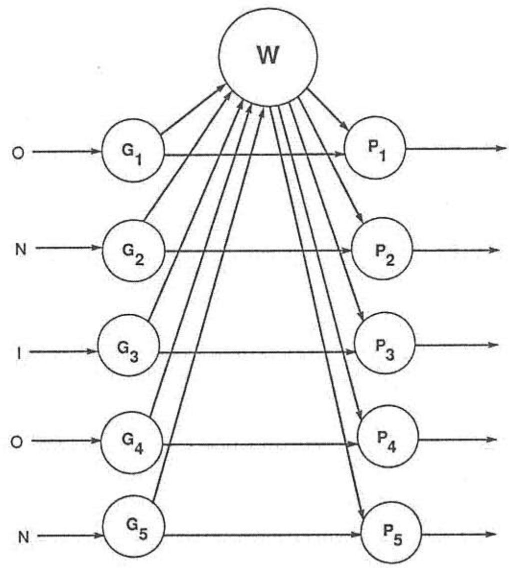

Figure 2: Overview of network structure when the five-grapheme sequence O N I O N is input to the print-to-sound model. The two *routes* in figure 1 are still evident. For each grapheme position, a corresponding set of graphemes and a set of phonemes exists, and each route is thus seen to be composed of five *paths*. There are no connections between two nodes in the same set anywhere in this network; all connections are forward-only.

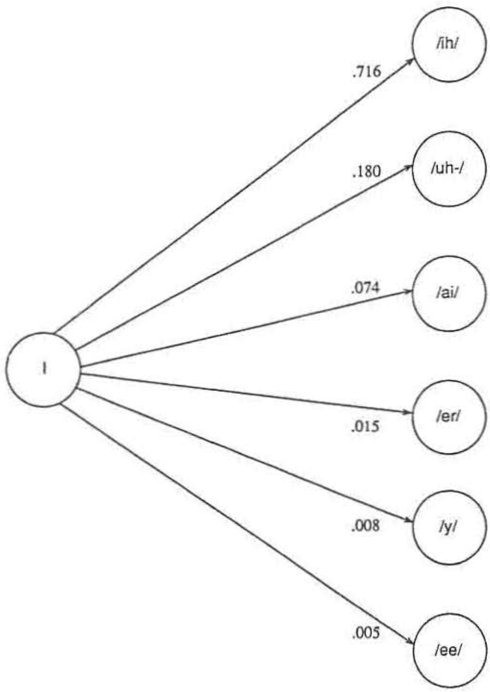

Figure 3: All connections from the grapheme node I in a specific position (path) to phoneme nodes in the same position. Each phoneme node receives additional incoming connections from other grapheme nodes, as well as connections from appropriate word nodes, that are not shown here. Each numeric weight represents the conditional probability that the grapheme I will be pronounced as the corresponding phoneme.

The Toronto Word Pool [21] was selected as the source for the model's lexicon because it contains a relatively large set (1080) of two-syllable words with a wide range of frequencies that were not preselected on the basis of orthographic or phonetic structure. Although this is only a fraction[2] of the words in the average person's lexicon, it provides a respectable set of word nodes comparing favorably in size to that used in many previous connectionist models developed by cognitive scientists [34]. In addition, the Toronto Word Pool is composed of two-syllable words. This word length was selected specifically in an attempt to move beyond the focus on monosyllabic words that has characterized previous models (e.g., [6,34]). The phonemic pronunciation of each word was taken from *Webster's New Collegiate Dictionary* (8th edition). Once the phonemes were determined for each word, the printed word was segmented into graphemes so that a single phoneme corresponded to a single grapheme.

Each grapheme node in the i^{th} grapheme set is connected to all word nodes in which that grapheme appeared in the i^{th} position. The weight on each link from a grapheme node to its word nodes is $1/n$, where n is the total number of words to which that grapheme connected. For example, the I in the third set G_3 of graphemes in figure 2 is connected to all word nodes with I in the third position, such as "onion," "union," "prison," and "amid." Since there are $n = 29$ such word nodes in set W, the weight on each of these links from I in the third position to word nodes is $1/n = .0345$.

Each word node also has forward connections to phoneme nodes in the appropriate position-specific phoneme sets. The word node for "onion," for example, has a connection to /y/ in the third set P_3 of phoneme nodes. With the exception of 17 words with common multiple pronunciations (e.g., the second phoneme of "content" is /ah/ or /uh-/ depending on whether "content" is a noun or adjective), a word node has one connection with a weight of 1.0 to a single phoneme node in each phoneme set (or no connections to some phoneme sets if the number of phoneme sets exceeded the number of phonemes in the word).

In summary, the associative network in the print-to-sound model consists of numerous positively weighted, forward connections forming multiple, position-specific paths through two routes (figure 2). From the perspective of an individual node in the network, it has multiple disjoint sets of connections with which to interact (figure 4). Each grapheme node g_i has an external input line and two sets of output connections going to word nodes and to phoneme nodes in the same position (see figure 4a). Each word node w_j has multiple inputs from graphemes and multiple outputs to phonemes that span the position-specific sets of graphemes and phonemes, respectively (see figure 4b). Finally, each phoneme receives inputs from two separate sets of connections: those from multiple word nodes, and those from multiple grapheme nodes in the same position (see figure 4c). There are no reverse

[2]Fifty words from the Toronto Word Pool were omitted from the simulations because they contained a silent letter other than H. These words could not be used without a modification of the Hanna, et al., correspondences. See [4, p. 5] for discussion.

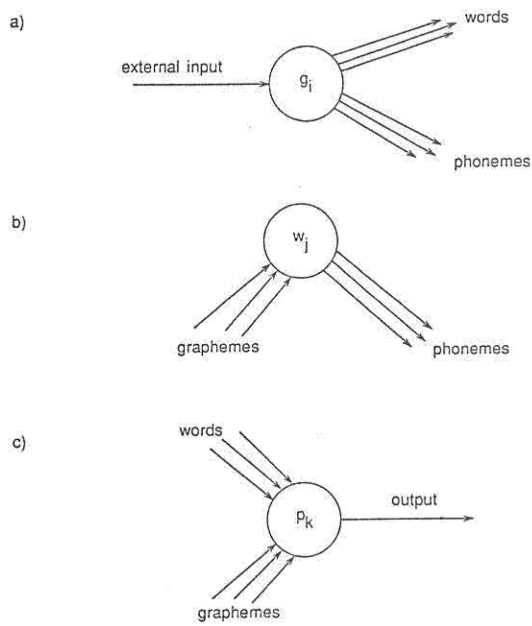

Figure 4: Summary of node connections: (a) Grapheme node g_i receives external input and sends output connections to both a set of possible phoneme nodes (GPC route) and a separate set of word nodes (lexical route); (b) Word node w_j receives input from grapheme nodes and sends output to phoneme nodes; (c) Phoneme node p_k receives one set of inputs from word nodes (lexical route) and a separate set of inputs from grapheme nodes (GPC route).

connections (phonemes to words or graphemes, words to graphemes), no inhibitory connections, and no connections between any two phonemes, any two graphemes, or any two words.

The complete network involved in a simulation is thus relatively large. For example, for the input graphemes O N I O N, there are a total of 2110 nodes (5 times 168 grapheme nodes, plus 5 times 48 phoneme nodes, plus 1030 word nodes) and roughly 12,000 forward connections. The use of a competitive activation mechanism (described in next section), however, avoids the need for more than a million lateral inhibitory connections that would normally be required to produce winner-take-all behavior among word nodes and among each position-specific set of phoneme nodes.

3. Competitive dynamics

Competitive interactions occur in many complex situations, and connectionist models are no exception. This section discusses the role of competition in connectionist models, and motivates and describes the concept of competitive activation mechanisms. The specific competitive activation mechanism used in the print-to-sound model is presented, and the sense in which this model represents an "interactive" dual-route hypothesis is explained.

3.1 Competition in connectionist models

To consider the issue of competition in connectionist models, it is useful at this point to explain some terminology. First, it is important to appreciate that the mechanism by which competition occurs in various complex systems may vary from situation to situation. *Direct (antagonistic) competition* is said to occur between two rivals A and B when A directly suppresses B's activities (e.g., wrestling match). *Indirect (allocational) competition* is said to occur when two rivals require and consume the same limited resource, the gain of one coming at the expense of the other (e.g., two animal populations competing for the same source of food). These two mechanisms for producing competition are not mutually exclusive.

In the following, it is assumed that each node in a connectionist model has a numeric activation level associated with it and an *activation mechanism*, a local algorithm that periodically updates the node's activation level as a function of input received from neighbor nodes. Two broad classes of connectionist models are distinguished: neural network models and associative network models. The term *neural network model* is used here to refer to connectionist models of neurophysiological systems (e.g., nodes represent neurons, links represent synapses, activation level represents neuron firing frequency, etc.) that typically adopt a distributed representation of concepts. The term *associative network model* is used to refer to "spreading activation" models developed in cognitive science and AI which have a local representation of concepts (e.g., nodes represent concepts, links represent relations between concepts, activation level represents probability/belief/desirability of concepts, etc.). Connectionist models involving semantic networks, and the non-semantic print-to-sound network described in this paper, are examples of associative networks.

A long-standing issue in the development of connectionist models has been how to integrate competitive and cooperative interactions between interconnected nodes so that meaningful model behavior emerges. When external activation is introduced into a network, some competitive influence on network dynamics is necessary to focus the spread of that activation and to avoid network saturation. In neural network models, direct/antagonistic competition has usually been used and implemented through negatively weighted inhibitory links between competing nodes (see figure 5a) [2,25]. Having "lateral" inhibitory links has proven very useful in neural network models, and is quite plausible in models of neurobiological circuitry given the overwhelming

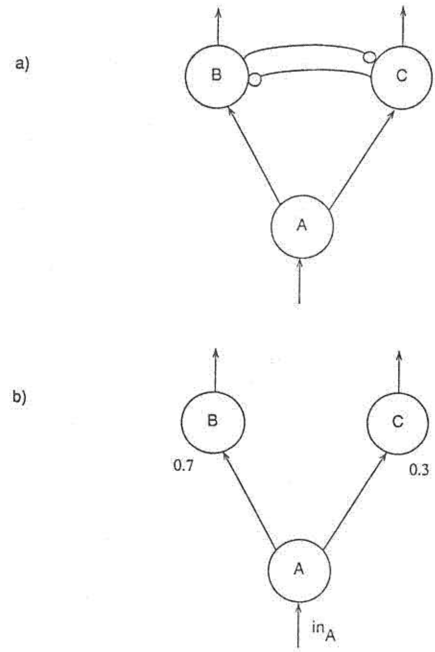

Figure 5: (a) Two directly competing nodes labeled B and C. Inhibitory connections, indicated by arcs with circular ends, are "lateral" (rather than "forward" or "backward") in that they are orthogonal to the flow of activation through this network fragment; (b) Two indirectly competing nodes B and C without lateral inhibitory links.

neurophysiological evidence of their importance.

Many associative network models developed in cognitive psychology and AI have relied on a similar approach by implementing competitive dynamics as direct competition. Lateral inhibitory connections have been widely used in associative network models, usually as a means for producing a *single-winner-takes-all* phenomenon [14,20,34,57]. This phenomenon is particularly relevant in models where there are a set of nodes which are conceived of as being mutually exclusive alternatives to one another. When these nodes get activated, it is desired that their initial, usually diffuse activation pattern be transformed into an equilibrium state in which one "winner" node is fully activated while all other competing nodes become fully inactive. In figure 5a, for example, if node B became highly activated for whatever reason, its inhibitory connection to node C would directly suppress node C's activation, which would also decrease inhibitory influences on node B from C, and typically lead to a stable equilibrium with node B as the sole winner.

While lateral inhibitory links provide a useful mechanism for implementing competitive interactions between nodes in associative networks, they also raise a number of important theoretical and practical issues. First, in contrast to neural network models where inhibitory links explicitly represent inhibitory synaptic connections, it is unclear exactly what inhibitory links represent in many associative network models of memory. Positively weighted links in associative networks generally represent measurable, application-specific associations between concepts. In contrast, negatively weighted inhibitory links are used to bring about application-independent functional properties (competitive interactions). There is no clear analog in contemporary psychological theories of associative memory for the inhibitory links that appear in connectionist models using associative networks.

This distinction can be clarified by considering associative networks in computational models that use the more traditional symbol-processing methods of cognitive science and AI rather than spreading activation. The networks in these symbol-processing models generally do *not* include inhibitory connections, and competitive interactions are imposed by an interpretative program with global access to the network (e.g., intersection search, or a generalization of it called parsimonious covering [42,43]). The point is that competitive interactions, traditionally viewed as a *functional* aspect of memory in symbol-processing models, are now being routinely implemented as a *structural* component of the network in connectionist models of memory. To our knowledge, little discussion of the implications of such a revision of network contents has occurred. For example, it is not immediately obvious in some associative networks exactly where inhibitory links should go, let alone how their weights should be assigned or measured [39].

Further, in many realistically sized networks of interest in cognitive science and AI, the number of inhibitory connections required to bring about the desired competitive interactions can be enormous. For example, if there are n nodes in a set and they must each directly inhibit one another to produce a single-winner-takes-all phenomenon during processing, then close to n^2 inhibitory connections would be required (e.g., only 1000 competing nodes would require almost 1,000,000 inhibitory connections). Thus, scaling up to large networks of the size often seen in cognitive science or AI models while using inhibitory connections in this fashion would clearly require a large number of connections and nodes with large fanouts. This is an important consideration not only when using simulated parallelism on a sequential (Von Neumann) machine, but also when connectionist models are implemented on parallel computers. The limited experience to date with actual implementation of connectionist models on parallel architecture hardware suggests that communication time between processors will be a major efficiency concern (e.g., [5]) and that in some situations speedup is adversely affected by large node fanouts [54].

Finally, in some important applications of associative networks a *multiple-winners-take-all* phenomenon rather than a single-winner-takes-all phenomenon is desired as the outcome of competitive interactions. In other words, there are situations where multiple winning nodes should be fully activated simultaneously while all other competing nodes should be fully inactive. General diagnostic problem-solving provides an example of this: a hypothesis or solution typically consists of one *or more* disorders which plausibly explain observed symptoms [42,43,38]. Implementing direct competition between nodes representing relevant disorders in such a situation is at best problematic. If a set of competing nodes with direct, mutually-inhibitory links between all nodes tries to sustain multiple winners simultaneously, these winners will tend to extinguish each other's activations. More selective use of inhibitory connections runs into the problem that, in some situations, two disorder nodes may be considered to be in competition, while in other situations the same two nodes may not be competitors and may actually "cooperate" to formulate a solution to a problem [39]. The competitive interrelationship between disorders during diagnostic inference is not simply a static, mutually-inhibitory relationship, but a more complex dynamic function of the network structure and the problem input (the latter being the set of present symptoms). Thus it is at least very difficult, if not impossible, to model these relationships through simple inhibitory links with static weights.

3.2 Competitive activation mechanisms

These difficulties with using inhibitory connections to implement competition in associative network models motivate an alternative approach. Rather than implementing direct or antagonistic competitive behavior through explicit structural features of a network (inhibitory links), indirect or allocational competition is introduced into the functional mechanism by which the spread of activation is controlled. When this is done, the resultant connectionst model is said to use a *competitive activation mechanism* [40,41].

In connectionist models using a competitive activation mechanism, as with many past network models, each node n_j transmits its activation level $a_j(t)$ at time t to neighboring nodes via weighted links, and these neighboring nodes update their own activation level based on activations received in this fashion. However, unlike previous models, with a competitive activation mechanism each such neighboring node n_m *actively* competes for the output from source node n_j. Further, the ability of a neighboring node n_m to compete for n_j's output increases as $a_m(t)$ increases.

Many formulations of a competitive activation mechanism are possible; one example follows. Let $a_k(t)$ and w_{ij} be restricted to the interval $[0, 1]$ for all i, j, and k, where w_{ij} is the connection strength *from* node n_j *to* node n_i. Note that in this example there are no inhibitory links since w_{ij} is always a positive number. Let the rate $\dot{a}_i(t)$ at which n_i's activation changes be given by

$$\dot{a}_i(t) = f_i(in_i(t), a_i(t)). \tag{3.1}$$

Here, f_i is a monotonically increasing function of $in_i(t)$, the total input activation to node n_i at time t resulting from external inputs and/or incoming connections from other nodes. For the competitive activation mechanisms discussed in this paper, calculation of $in_i(t)$ for each node n_i is arranged so that $0 \leq in_i(t) \leq 1$. This is achieved using

$$in_i(t) = 1 - \Pi_k(1 - \text{out}_{ik}(t)) \tag{3.2}$$

where $\text{out}_{ik}(t)$ is the output from node n_k arriving at node n_i at time t, or is an external input to n_i. As long as each out_{ik} value satisfies $0 \leq \text{out}_{ik} \leq 1$, equation (3.2) guarantees that the resultant in_i is also restricted to this range. Equation (3.2) can be viewed as combining individual inputs out_{ik} in a nonlinear, accumulative fashion. If each and every individual out_{ik} is zero, then so is in_i. If any out_{ik} is nonzero, then in_i is non-zero, and in general, the more individual out_{ik} values that are non-zero, the greater the resultant in_i. According to Eq. (2), in_i is a monotonically increasing function of every out_{ik}, and can be viewed as a numerical version of a logical OR operation.

The activation mechanism described so far as equations (3.1) and (3.2) does not differ in any fundamental way from many non-competitive activation mechanisms. To introduce allocational competition into this model, consider the perspective of a node n_k computing out_{ik}, its output to node n_i. Let

$$\text{out}_{ik}(t) = c_{ik}(t) \cdot a_k(t) \tag{3.3}$$

where $c_{ik}(t)$ is the competitive strength of node n_i determining how much of $a_k(t)$ reaches n_i. As a specific example, let

$$c_{ik}(t) = \frac{w_{ik} \cdot a_i(t)}{\sum_m w_{mk} \cdot a_m(t)} \tag{3.4}$$

where m ranges over nodes to which n_k sends connections. If the denominator in equation (3.4) is zero then the numerator is also zero, and by definition we let $c_{ik}(t) = w_{ik}$ in this case. Note that $0 \leq c_{ik} \leq 1$, so by equation (3.4), $0 \leq \text{out}_{ik} \leq 1$, and thus the total input in_i received by any node n_i is guaranteed to satisfy $0 \leq \text{in}_i \leq 1$ (see equation (3.2) and discussion following it). Further, it follows from equation (3.3) and (3.4) that the total output of any node n_k is $\sum_m \text{out}_{mk} = a_k(t)$.

The key point here is the appearance of $a_i(t)$ and $a_m(t)$, the current activation levels of nodes *receiving* activation from node n_k, in the formula for c_{ik}. This is what makes this a competitive activation mechanism involving allocational competition. Node n_i "competes" for n_k's activation such that the portion of a_k it receives increases as a_i increases. Conversely, if some competitor n_m of n_i receives input from n_k, then by equation (3.3) and (3.4) the amount of input that n_i receives from n_k will decrease as a_m increases. This can be contrasted with the situation where a *non*-competitive activation mechanism is used. Typically, with a non-competitive activation mechanism, $c_{ik}(t) = w_{ik}$, a constant value for all time in non-adaptive networks, so $\text{out}_{ik}(t) = w_{ik}a_k(t)$. In this case the fraction of $a_k(t)$ distributed to node n_i remains constant with time and a_k is not allocated competitively to nodes n_m to which n_k connects.

Iteration	A	B	C
0	0	0	0
5	.344	.022	.007
10	.613	.216	.013
20	.865	.596	.011
40	.984	.937	.004
60	.998	.992	.001
100	1.000	1.000	.000

Table 1: Node activations for the network pictured in figure 5b using the competitive activation mechanism described by equations (3.2–3.5). Derived numerically using time quantization of 0.1 unit of time per iteration.

As a simple specific example, consider the 3-node network in Figure 5b where all nodes start with zero activation, and a constant external input of 1.0 is applied to node A starting at $t = 0$. Each node uses the competitive activation mechanism described by equations (3.1–3.4) where, in equation (3.1, we let

$$\dot{a}_i = [\text{in}_i - a_i(1 - \text{in}_i)] \cdot (1 - a_i) \tag{3.5}$$

The first factor here, $\text{in}_i - a_i\,(1 - a_i)$, can be positive or negative, and ranges from -1 to $+1$. This factor causes a_i to increase whenever $\text{in}_i > a_i\,(1 - \text{in}_i)$, and to decrease whenever $\text{in}_i < a_i(1 - \text{in}_i)$, and insures that $a_i \geq 0$. The second factor, $1 - a_i$, insures that a_i has 1.0 as a maximum value. For this specific example, approximate activation of nodes with time is given in table 1. As the external input activates node A, nodes B and C are both initially partially activated, but as equilibrium is approached a winner-takes-all phenomenon appears (node B fully activated, node C fully inactivated). Although no inhibitory links exist between nodes B and C, an indirect inhibitory interaction ("virtual lateral inhibition") between these nodes is apparent as a result of the allocational competition controlling how node A distributes activation to nodes B and C.

The behavior of connectionist models using a competitive activation mechanism has been studied so far primarily through small scale simulations [41,39] and limited theoretical analysis of simple networks [51]. This work has clearly demonstrated many useful properties of competitive activation mechanisms: circumscribed network activation, trajectories leading to an equilibrium point (attractor), ability of suitable formulations of a competitive activation mechanism to produce winner-takes-all behavior in the absence of inhibitory links, and context-sensitivity of the winner-takes-all phenomenon. These results suggest that allocational competition can implement the types of competitive interactions needed in many connectionist models without the problems associated with lateral inhibitory links as outlined earlier (theoretical representational issues, assignment of inhibitory weights, large number of connections needed, etc.). However, this previous work with competitive activation mechanisms has only examined small, uncomplicated networks. Not uncommonly, methods developed for use in connectionist models in the past have been successfully applied to small, simple problems, only to have subsequent study show that they do not scale up well to larger and more realistic situations. The model of print-to-sound transformation described in the next section begins to address this issue for competitive activation mechanisms. It demonstrates for the first time that connectionist models using competitive activation mechanisms can be developed for networks of a size and complexity found in typical cognitive science and AI applications.

3.3 Activation rule for the print-to-sound connectionist model

We now consider the specific competitive activation rule used in the print-to-sound model. All nodes in the associative network follow a similar "rule" in updating their activation and output levels during a simulation. This rule involves strictly local computations as determined by equations (3.1–3.4) after adjustment of these equations to accommodate the complexities arising from having multiple routes and paths through a network composed of multiple classes of nodes (grapheme, word, and phoneme nodes). The exact form of this rule (described below) was determined initially in an intuitive fashion, and then modified based on preliminary simulations using a small, abstract network. The purpose of these initial exploratory simulations with a prototype network was to produce a specific competitive activation mechanism that provided clear-cut winner-takes-all behavior (i.e., correct "winner" nodes that had activation above .99 and all "loser" nodes with activation below 0.01). The network used in this preliminary work had only 16 word nodes with an average of about 3 graphemes/phonemes per word. There were 4 possible graphemes and 4 possible phonemes per position in a word, and 9 arbitrarily-weighted connections between these graphemes and phonemes. The activation rule described below is the best of the limited number of variations examined during this exploratory work. In the following, "preliminary simulations" refer to simulations done with this small network.

Starting with $a_i(0) = 0.0$, each node in the print-to-sound model uses the following specific version of equation (3.1) to update its activation level:

$$\dot{a}_i = k_i \cdot [\mathrm{in}_i - 2a_i(1 - \mathrm{in}_i)](1 - a_i) \tag{3.6}$$

This is the same as equation (3.5) except that two constants have been introduced. The value $2a_i(1 - \mathrm{in}_i)$ in the second factor is used rather than $a_i(1 - \mathrm{in}_i)$ as in equation (3.5) because in the preliminary simulations this alteration was observed to result in much cleaner winner-takes-all behavior. The other new constant k_i was 1.0 for phoneme and grapheme nodes. For word nodes, the value of k_i was a logarithmic function of p_i, a node's prior probability.[3] We introduced k_i to allow analysis of word frequency effects.

[3]Specifically, k_i is an increasing function of p_i, the prior probability of the i^{th} word, given by $k_i = .45\log(p_i \cdot 10^6)$. The prior probabilities of words used in the print-to-sound lexicon ranged from about 1×10^{-6} to 1815×10^{-6}. The natural logarithm function is

Input to all nodes is based on equation (3.2), but the details of how this was achieved differ depending on whether the node involved is a grapheme, word or phoneme node. In the particularly simple case of a grapheme node g_i (figure 4a), only a single "external input," designated out_{io}, is present, and equation (3.2) simplifies to $\mathrm{in}_i = \mathrm{out}_{io}$. In the simulations the value of in_i for grapheme g_i indicates the presence ($\mathrm{in}_i = 1.0$) or absence ($\mathrm{in}_i = 0.0$) of the grapheme represented by node g_i. In this special case, with a constant $\mathrm{in}_i = 1.0$, the activation rule (equation (3.6)) simplifies to $\dot{a}_i = 1 - a_i$ having solution $a_i = 1 - e^{-t}$ for $t \geq 0$. Thus a_i initially increases rapidly then progressively more slowly as a_i asymptotically approaches 1.0. Input to a grapheme node is thus so simple that it directly indicates whether the node should turn on or off.

For word and phoneme nodes the situation is more complex because these nodes receive inputs along multiple paths (figures 4b, 4c). The input along any single path is again determined using equation (3.2), but inputs along different paths may provide conflicting information about whether or not the receiving word or phoneme node should be activated. Further, input along each path changes continually and sometimes dramatically during a simulation as nodes on that path compete for available output from nodes sending them activation. It is therefore useful to compute the total input in_i to a word or phoneme node using an *input combining function* of the individual inputs as determined for each individual path using equation (3.2). A similar approach has been used by others, such as with "conjunctive" or "sigma-pi units" [47,20, p. 73]

For word nodes, each node w_i receives inputs via n paths (figure 4b) where n is the number of grapheme/phoneme positions in the word represented by node w_i. Each path-specific input in_{ip} is determined by equation (3.2): $\mathrm{in}_{ip} = 1 - \Pi_k(1 - \mathrm{out}_{ik})$. Initially, we combined these individual path-specific inputs in_{ip} by taking the resultant in_i in equation (3.6) to be the average of the n path-specific inputs. This insured that each word node receiving any input at all would become at least partially activated. Limited exploratory simulations with the small, abstract network described earlier indicated that using a different input-combining function (product rather than average) for the second occurrence of in_i in equation (3.6) produced cleaner winner-takes-all behavior among word nodes, so this latter function was used in the print-to-sound model. In retrospect, this improved performance makes sense because it results in a larger subtrahend in the first factor in equation (3.6) when any grapheme in a word is missing, thereby lowering that word node's activation and hence its ability to compete.

For phoneme nodes, each node p_i receives inputs via two separate position-specific paths (figure 4c). These are designated in_{iW} for the path in the lexical route, and in_{iG} for the path in the GPC route. Both in_{iW} and

---

used because of the large ratio between these two endpoints; multiplication of p_i by 10^6 is used so $k_i > 0$ for all words, and scaling by .45 makes k_i values for low frequency words lie around 1.0, the value of k_i used for grapheme and phoneme nodes, and k_i for high frequency words lie between 2.0 and 3.5.

in_{iG} are calculated using equation (3.2), and can differ from one another greatly in value (e.g., in irregularly spelled words). We wished to resolve such conflicting inputs in a symmetric fashion such that a phoneme node did not have to be concerned about the class of nodes (words or graphemes) responsible for input along a path (i.e., substituting in_{iG} and in_{iW} for each other in the input-combining function should result in the same function). Initially, we combined these path-specific inputs using simply

$$\text{in}_{i\text{AND}} = \text{in}_{iW} \cdot \text{in}_{iG}, \tag{3.7}$$

reasoning that both the lexical routes *and* the GPC route connections should be active in order for phoneme p_i to be activated. The product in equation (3.7) can be viewed as a numerical version of a logical AND operation. While this works reasonably well when input graphemes form a word, the preliminary simulations revealed that if input graphemes are non-words then clear-cut winner-takes-all phonemes usually did not occur. The reason for this is that word nodes only partially match the input graphemes of a non-word and thus are weakly activated, so in_{iW} is usually small. While a large lexicon having more partially-activated word nodes during a simulation could improve this situation, we elected instead to use a more complex input-combining function. The revised input-combining function first determines resultant input to be proportional to the extent that input is arriving at p_i via both routes (using $\text{in}_{i\text{AND}}$), but as p_i's activation level increases, the resultant input gradually shifts to become proportional to the extent that either input route is active (using $\text{in}_{i\text{OR}}$). This is done as follows. Recalling that equation (3.2) can be viewed as a numerical version of a logical OR operation, analogously let

$$\text{in}_{i\text{OR}} = 1 - \Pi_R (1 - \text{in}_{iR}) \tag{3.8}$$

where R is W or G.[4] The input combining function for phonemes as used in the full print-to-sound model is then

$$\text{in}_i = \text{in}_{i\text{AND}} (1 - a_i) + \text{in}_{i\text{OR}}\, a_i. \tag{3.9}$$

Note that this behaves precisely as described above. Initially a_i, the activation of the i^{th} phoneme node p_i, is very small, so $\text{in}_i \cong \text{in}_{i\text{AND}}$. Subsequently, as p_i becomes more activated, in_i gradually shifts progressively closer to being $\text{in}_{i\text{OR}}$.

Finally, for all nodes in the network the output out_{ik} from node n_k to node n_i is determined by equations (3.3) and (3.4). This is done separately for each path to which a node sends output. Each grapheme node g sends output via connections in two paths (in the lexical and GPC routes; see figure 4a). Output from g is divided up competitively according to equations (3.3) and (3.4) among the connections to word nodes, and *separately* is divided up

[4]It follows that $\text{in}_{i0R} = 1 - \Pi_k\ (1 - \text{out}_{ik})$, where k ranges over *all* input connections to p_i from *both* graphemes and words.

competitively among the connections to phoneme nodes. Stated otherwise, phoneme nodes do not directly compete with word nodes, nor vice versa, in extracting their share of g's output, so the two routes do not compete against each other for g's output. A similar situation holds for each of the position-specific path outputs from a word node to phonemes in the sense that phoneme nodes in one path (e.g., those in $P3$ in figure 3) do not compete against phoneme nodes in other paths (e.g., $P4$).

3.4 Independent versus interactive routes

We are now in a position to clarify the sense in which the two routes in the print-to-sound model are interactive. As described above, an individual activated grapheme g effectively distributes two identical "copies" of its activation along its outputs, one to phoneme nodes and one to word nodes. From g's local perspective, events along the GPC route do not affect the total amount or distribution of activation being sent along the lexical route, and vice versa. Further, there is no direct or "lateral" influence of the GPC and lexical routes upon one another. Thus, if in addition no retrograde influences were present (i.e., influences flowing from phonemes back through the network), the two routes in the print-to-sound model would represent an *independent* dual-route theory of information processing.

The most common way that retrograde influences are implemented in connectionist models is through "backward" connections over which activation flows in a reverse direction. For example, such reverse connections were used as a critical aspect of the interactive activation model of letter perception in context [34,46]. If reverse connections were present in the print-to-sound network described in this paper, then the two routes involved could be characterized as representing a *directly interactive* dual-route theory. The qualifier "directly interactive" is used in the sense that such connections would permit, for example, phoneme nodes to increase or suppress directly the activation levels of word nodes or grapheme nodes. In such a situation, the activation of phoneme nodes by grapheme nodes via the GPC route could directly exert an influence on the activation of word nodes in the lexical route, making the two routes strongly interactive.

The print-to-sound model described in this paper involves two routes of information flow which are neither independent nor directly interactive. They are not directly interactive in that no reverse connections exist that permit phoneme nodes to directly influence activation levels of word or grapheme nodes. However, neither are the two routes completely independent. Although there is no reverse flow of activation, there is limited reverse flow of information that steers and focuses the foward flow of activation. Nodes distributing their activation in a forward direction are influenced in the manner in which their activation is parcelled out to receiving nodes by the activation levels of those receiving nodes (see discussion of competitive activation mechanisms above). Activation of a phoneme node via, for example, the lexical route can influence how much activation that phoneme node receives from a grapheme node via the GPC route, but does not influence the activation of the grapheme node itself. We thus say that the network of the print-to-sound model represents an *indirectly interactive* dual-route model. The two routes do not compete against each other for grapheme node activation, but each route does influence how the other route's competition is resolved.

4. Model performance

The print-to-sound connectionist model involves a large associative network based on the best published empirical data known to the authors [26,21,56]. Nevertheless, the scope of the network is significantly circumscribed in obvious ways: no semantic component is present, the lexicon is limited in size, and no word-specific morphological, syllabic or segmentation information is used. The print-to-sound model should thus be viewed as only a first approximation to print-to-sound transformation, and its evaluation necessarily could have only limited goals. One goal was to establish whether a competitive activation mechanism could produce suitable winner-takes-all behavior among word nodes and phoneme nodes in the absence of lateral inhibitory connections. As explained earlier, previous work with competitive activation mechanisms has used only relatively trivial networks. It was not obvious *a priori* that this approach could be scaled up to the complex print-to-sound network involving conflicting input signals to phoneme nodes arriving over two separate routes.

The second goal was to examine qualitatively how the performance of an indirectly interactive dual-route model would correlate with published data on reading aloud single words and non-words. As discussed earlier, skilled adult subjects read aloud regular words faster than exception ("irregular") words in the lower frequency range, and read all types of words faster than non-words. The issue here is whether or not a dual route model is at least consistent with these findings and other relevant performance data.

In the following, we briefly summarize word regularity as it relates to the print-to-sound model, and give an example of the model's performance during a single simulation. Following this the results of limited but systematic simulations with the intact and "damaged" model are presented.

4.1 Regularity metric

The notion of word regularity or irregularity, or of "exception words," is complex and controversial. The traditional notion of regularity assumes a single preferred pronunciation for letters and multi-letter graphemes. A set of rules can be formulated that captures reasonably well the normal correspondences between letters and sounds [55]. Recently, however, the idea of rule-based grapheme-to-phoneme correspondences has been challenged by the notion that pronunciation may be based on associations between sounds and print segments of various sizes that are probabilistically derived from known words [37]. Values for rule-based correspondences and for word-based correspondences often coincide, and there is some evidence that words in which they do coincide are read most easily and quickly [23]. Nonetheless, these two types of "regularity" can be manipulated independently when the size of the segment to be pronounced is larger than a single grapheme [31]. Various arguments have been advanced favoring both types of "regularity" as the determinant of ease of pronunciation and there has even been a proposal that pronunciation latencies reflect an interaction of the two [45]. Most of these arguments are based on the reading of monosyllabic words; two-syllable words could only be expected to involve even more complexity.

Given the difficulties involved in characterizing word regularity, a simple *regularity metric* was adopted for this study based on the frequency of grapheme-to-phoneme correspondences in a word. This metric is a compromise between the two positions summarized above in that it maintains the grapheme as the unit size (as does the rule-based position) but it computes probabilities of correspondence based on the number of words in which a particular correspondence occurs (as does the "word-based" position.) The intent was not to dictate what "regularity" should be, but to provide an objective quantitative estimate of the strength of particular correspondences for unbiased comparison with simulation results.

Perhaps the simplest word regularity metric of this sort would be the average of the grapheme-phoneme correspondence (GPC) frequencies occurring in a given word. However, such a metric ignores the fact that many English words considered to be "irregular" have a single very uncommon vowel correspondence (e.g., "many"). Thus, the metric described here weights low-frequency GPC's more heavily in forming the average. Let prob_i be the relative frequency with which the grapheme-to-phoneme correspondence in the i^{th} position in a word occurs in English [4]. Then the regularity of the k^{th} word, designated R_k, is given by

$$R_k = \frac{\sum_{i=1}^{n}(1.05 - \text{prob}_i)\text{prob}_i}{\sum_{i=1}^{n}(1.05 - \text{prob}_i)} \tag{4.1}$$

where n is the number of grapheme-phoneme positions in the word. In averaging the individual prob_i values, this formula weights each by $1.05 - \text{prob}_i$, thus giving lower frequency prob_i's a higher weight. The value 1.05 was used rather than 1.0 so that all GPC frequencies are counted, albeit slightly, in this average, even those with $\text{prob}_i = 1.0$.

Applying the word regularity metric R_k described by equation (4.1) to the model's lexicon of 1030 words from the Toronto Word Pool, *prior to running the simulations whose results are described below*, convinced the authors that this metric provided a coarse but reasonable measure of word regularity. For example, consider the three words in table 2. The highly regular word "needle" has $R_k = 0.99$, the less regular word "disturb" has $R_k = 0.41$, and the very irregular word "onion" has $R_k = 0.12$.

It should be noted that the regularity metric R_k is based on an input string that is already parsed into letter segments (graphemes). Thus it does not capture possible irregularities that might result from a mis-parsed letter

	"needle"		"disturb"		"onion"	
i	GCP_i	$Prob_i$	GPC_i	$Prob_i$	GCP_i	$Prob_i$
1	N → n	1.0	$D \rightarrow d$	1.0	O → uh+	0.007
2	EE → ee	0.98	I → ih	0.72	N → n	0.967
3	D → d	1.0	S → s	0.87	I → y	0.008
4	LE → ul	1.0	T → t	0.97	O → uh-	0.269
5			U → er	0.08	N → n	0.975
6			R → r	1.0		
7			B → b	1.0		

Table 2: Examples of words and the frequencies of their grapheme-phoneme correspondences to illustrate the word regularity metric R_k.

string. For example, "design" might be thought of as irregular because of the silent G; in the system developed here (following [26]) GN is considered to be a grapheme with high probability (=1.0) of pronunciation as /n/. Since the print-to-sound model discussed in this paper starts with a segmented input string, a regularity metric based solely on the probability of grapheme-to-phoneme correspondences was deemed the most appropriate one to use.

4.2 An example simulation

The print-to-sound model was implemented using MIRRORS, a general purpose software environment for developing connectionist models [16,17]. All simulations were run on a single-processor DEC MicroVAX/2 under Unix© using single-precision arithmetic. A time step of 0.1 was used during numeric calculations. Input characters were manually grouped into graphemes.

A single brief example of a representative simulation with the print-to-sound model is given here to illustrate its ability to activate correctly word and phoneme nodes with a very sharply defined winner-takes-all performance. A sequence of five graphemes, O N I O N, representing an irregular word ($R_k = 0.12$), serves as input to the model. The overall network structure for this input has already been seen (figure 2). Table 3 gives activation levels as a function of time for selected nodes in the network. The symbol "-" means "inactive" ($a_i < .001$) and the symbol "*****" means "saturated" ($a_i > .99$). Each 10 iterations (first column in table 3) represents one unit of simulated time.

Grapheme nodes quickly become saturated (second column). While a fair number of word nodes are activated early in the simulation, activations for only two of these word nodes are given here (columns 3–4). These nodes represent the target word "onion" and one of its orthographic neighbors, "union," which has four of five graphemes in common with the target word. These are the most highly activated word nodes during this simulation. Early in the simulation, the node representing the orthographic neighbor "union" is *more* activated than the node for "onion." However, eventually "onion"

Iterations	Grapheme	Words		Phonemes				
		"onion"	"union"	$/uh\text{-}/_1$	$/uh+/_1$	$/ih/_3$	$/y/_3$	$/n/_5$
0	-	-	-	-	-	-	-	-
10	.651	.063	.134	.001	-	.001	-	.029
20	.878	.211	.164	.002	-	.001	-	.166
30	.958	.383	.097	.001	-	-	-	.360
40	****	.657	.023	-	-	-	.253	.613
50	****	.884	.002	-	.272	-	.658	.803
60	****	.967	-	-	.722	-	.875	.942
70	****	****	-	-	.879	-	.956	.979
80	****	****	-	-	.966	-	.985	****
90	****	****	-	-	****	-	****	****

Table 3: Activation of selected nodes in the print-to-sound network following input of the grapheme sequence O N I O N starting at $t = 0$. Subscripts of phonemes indicate their positions. The entry "-" indicates a node is inactive ($a_i < .001$) and the entry "****" indicates a node is fully active ($a_i > .99$).

dominates and becomes fully activated, while activation of all other word nodes (both those shown here and all others) dies out. The clear-cut winner-takes-all behavior arises completely through allocational competition, with "onion" eventually dominating because of its perfect match with the input graphemes in this case. The larger early activation of the node representing the orthographic neighbor "union" arises primarily as a word frequency effect. The prior probability of "union" ($p_i = 182 \times 10^{-6}$; $k_i = 2.34$) is much larger than that of "onion" ($p_i = 15 \times 10^{-6}$; $k_i = 1.22$). Thus, even though the "union" node does not initially receive as much input from graphemes on all paths as the "onion" node does, it activates more quickly initially due to the larger k_i in equation (3.6).

The last three columns in table 3 illustrate activations of selected phoneme nodes. Allocational competition results in clear-cut, winner-takes-all activation of exactly those phoneme nodes representing the correct pronunciation of the word "onion." Activation of phoneme nodes is slower than that of word nodes because, especially early on, their activation depends on receiving significant input from *both* the lexical and GPC routes (equations (3.7,3.9)). The rightmost column illustrates the mapping of grapheme N_5 to phoneme $/n/_5$ in the fifth position. In this case, the GPC route connection $N_5 \rightarrow /n/_5$ in the final position has the large weight .975 (table 2). Further, the most highly active word nodes ("onion" and "union") both have a $/n/_5$ in their phonemic realization. Thus, $/n/_5$ receives reinforcing input simultaneously from the lexical and GPC routes, and rapidly activates with little significant competition. In contrast, in the third position grapheme I_3 connects to six phonemes (figure 3). Activation levels for two of these phonemes, $/ih/_3$ and $/y/_3$, are given in table 3. Weights on the GPC route connections are .716 for $I_3 \rightarrow /ih/_3$ and .008 for $I_3 \rightarrow /y/_3$, the latter being the correct phonemic realization of I_3 in "onion." The very low weight on the GPC connection to $/y/_3$ and support for $/ih/_3$ from the word route (e.g., from "prison," "exist," etc., in this case) result in the slow activation of $/y/_3$ relative to $/n/_5$; see table 3. However, a clean winner-takes-all activation of $/y/_3$ still eventually occurs. A similar situation holds for $/uh\text{-}/_1$ and $/uh+/_1$ (table 3).

4.3 Simulations with intact model

To examine word frequency and regularity effects on the model's performance, data consisting of four sets of words with 16 words per set were used to test the model. Words for each of these sets were selected from those in the associative network's lexicon based on word frequency (high vs. low) and regularity (very regular vs. very irregular). One test set consisted of only high-frequency regular words; another of only high-frequency irregular words; a third of only low-frequency regular words, and the fourth of low-frequency irregular words. Word frequency was obtained from the Kucera and Francis [33] norms given in the Toronto Word Pool [21]. Word frequency ranged from 1×10^{-6} to about 1800×10^{-6} with a median of 30×10^{-6} and a mean of 77×10^{-6} for the 1030 words in the model's lexicon. "High-frequency" words were arbitrarily defined as those having frequencies in the top quarter of all frequencies ($> 80 \times 10^{-6}$) and "low-frequency" words as those with frequencies in the bottom quarter of all frequencies ($< 18 \times 10^{-6}$). High- and low-frequency regular words were selected by starting at the top of a regularity ranking of the 1030 words based on R_k and systematically selecting 16 words which met the word frequency criteria stated above. In a similar fashion, high and low-frequency irregular words were selected by starting at the bottom of the regularity ranking of the 1030 words.

The range of word frequencies for the set of low-frequency, irregular words was 3×10^{-6} to 15×10^{-6} with a mean of 10.2×10^{-6}. The set of low-frequency, regular words had similar values with word frequencies ranging from 4×10^{-6} to 18×10^{-6} and a mean of 10.8×10^{-6}. The set of high-frequency, irregular words had word frequencies of 114×10^{-6} to 1236×10^{-6} with a mean of 377×10^{-6}, while the set of high-frequency regular words had a frequency range of 94×10^{-6} to 831×10^{-6} and a mean of 190×10^{-6}. To achieve more similar means between the two sets of high-frequency words, the four words with the highest frequencies in the irregular word set were omitted and replaced with four newly selected irregular words which were obtained by continuing up the lexicon list (from the bottom of the regularity ranking based on R_k) and selecting the next four words with frequencies exceeding 80×10^{-6}. The new mean for this set of 16 words was then 189×10^{-6}, which was very similar to the mean frequency for the high-frequency, regular words of 190×10^{-6}. The four sets of test words which were used to evaluate the print-to-sound model are listed in the appendix.

In all 64 runs where a word contained in the model's lexicon was introduced as input, the correct set of phonemes eventually attained 1.0 activation, and all remaining phoneme nodes in each of the phoneme sets had 0.0

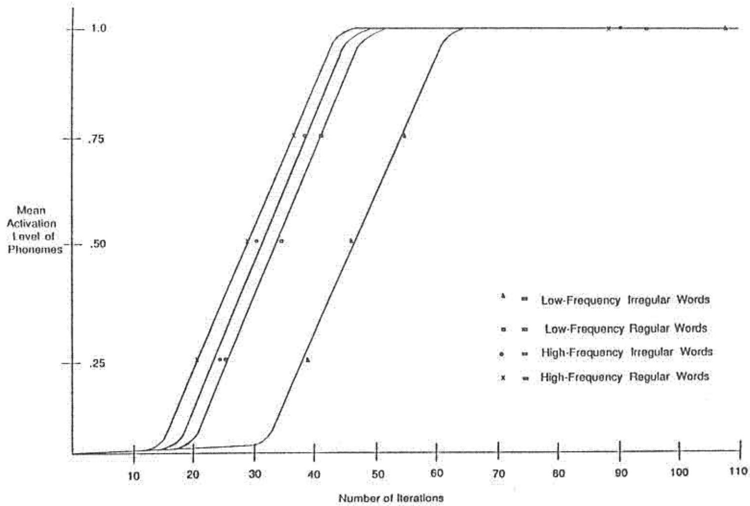

Figure 6: Mean activation of phoneme nodes for the four classes of test words as a function of time.

activations. Thus, clear-cut winner-takes-all behavior for the correct phoneme node always occurred with these simulations in a fashion similar to that demonstrated with "onion" above (table 3).

Measurements were recorded with the print-to-sound model of the time required for the correct set of phoneme nodes in each of the 64 words to attain 0.25, 0.50, 0.75 and 0.999 activation. These data were averaged for the 16 words in each data set. Figure 6, based on these averages, shows clearly that phonemes for low-frequency irregular words had a delayed activation rate compared to phonemes in the other three sets of words. Phoneme activation levels of "winning" nodes followed a sigmoid curve in asymptotically approaching a value of 1.0.

Let variable $t_{.5}$ represent the time (iterations) required for a word's phoneme nodes or a word node to attain a 0.50 activation level. Variable $t_{.5}$ provides a representative value (see figure 6) or "time constant" for statistical analysis of the differences in rate of phoneme activation. Averaged $t_{.5}$ values are presented in table 4 (columns 3–4) for each of the four sets of test words where k_i is based on word frequency.

A one-way analysis of variance with two factors (regularity and frequency) assessed by contrast was used to evaluate the timing data for the four sets of words in table 4. To compensate for a rather large variation in standard deviations among the data sets, log time was used as the dependent variable. The main effects of regularity (F 3,60 $= 48.96$, $p < 0.0001$), frequency (F 3,60 $= 106.46$, $p < 0.0001$) and interaction of regularity and frequency (F 3,60 $= 17.70$, $p < 0.0001$) were all highly significant. Based on the Student

Word Group	n	Using Word Frequencies		Ignoring Word Frequencies	
		phonemes $t_{.5} \pm sd$	word $t_{.5} \pm sd$	phonemes $t_{.5} \pm sd$	word $t_{.5} \pm sd$
1. High frequency Regular	16	27.9 ± 0.73	18.8 ± 1.29	32.5 ± 1.08	25.1 ± 1.61
2. High frequency Irregular	16	30.6 ± 2.39	18.9 ± .89	38.7 ± 5.00	26.8 ± 3.10
3. Low frequency Regular	16	33.3 ± 2.50	26.1 ± 4.06	31.6 ± 0.74	23.8 ± 1.53
4. Low frequency Irregular	16	46.7 ± 9.45	35.3 ± 14.27	40.3 ± 4.59	25.7 ± 2.91
5. Non-words	10	66.6 ± 12.55	NA	68.6 ± 11.79	NA

Table 4: Mean times for phoneme and word nodes to attain 0.50 activation in five different word groups with two different methods of spreading activation. Columns labeled "word" refer to the word node representing the external input of a sequence of graphemes.

Newman Keuls test (a post hoc multiple comparison test), no significant difference was found in the mean $t_{.5}$ values of phonemes for high-frequency regular versus high-frequency irregular words; however, the mean $t_{.5}$ values of phonemes for low-frequency regular words versus low-frequency irregular words were significantly different ($p < 0.005$). Thus, the interaction effect of regularity and frequency is quite evident in that high-frequency regular and irregular words had very similar phoneme activation times (27.9 and 30.6, respectively), whereas low-frequency regular words had significantly faster phoneme activation times than low-frequency irregular words (33.3 versus 46.7).

The mean word $t_{.5}$ value for high-frequency regular words (18.8 iterations) was very similar to the mean word $t_{.5}$ value for high-frequency irregular words (18.9 iterations). Unexpectedly, the mean word $t_{.5}$ value for low-frequency regular words (26.1 iterations) was significantly less ($p < .005$) than that of low-frequency irregular words (35.3 iterations). It had been anticipated that these latter two values would be approximately the same since only word frequency and activated word nodes representing competing orthographic neighbors (and *not* regularity) should influence the rate of word node activation in the model. To establish whether this unanticipated finding with word activations was responsible for the results described in the preceding paragraph, the four low-frequency irregular words with the largest word $t_{.5}$ values were omitted from the original set of 16 low-frequency irregular words to create a modified test set. The new mean word $t_{.5}$ value for this modified test set containing 12 words from the original low-frequency irregular word set was 27.8 iterations, which is very similar to the mean word $t_{.5}$ value of 26.1 iterations for the set of 16 low-frequency regular words. Repeating the simulations with the modified test set, the mean $t_{.5}$ value for phonemes was 42.7 ± 6.27. This is still significantly greater than the mean phoneme $t_{.5}$ value of 33.3 ± 2.5 for the low-frequency regular word set. This is an important point: it indicates that the slower activation times for phoneme nodes of low-frequency irregular words versus regular words was not primarily the result of differences in rates of word activations, but rather of differences in regularity of the words in the data set.

Three of the four "outlier" words in the low-frequency irregular word set (which had word $t_{.5}$ values of 45, 54, 56 and 70 iterations while the mean for the remaining 12 words was 27.8 iterations) had very low frequencies (3×10^{-6} and 4×10^{-6}) and had orthographic neighbors with very high frequencies; thus, the rate of word node activation for these particular words was significantly slower than for the remaining words in the data set. The fourth outlier word, "resort," had a moderately low frequency (12×10^{-6}) but its orthographic neighbor "report" was a very significant competitor (because it shared 5 of 6 graphemes with "resort" and had a high frequency of 174×10^{-6}), so the rate of word node activation for "resort" was quite slow. In the relatively small set of 16 words, the low-frequency regular words did not happen to have any distinct outliers. If a larger lexicon was used and the data sets contained more words, the mean word $t_{.5}$ values for these two data sets would probably be very similar.

To confirm that the phoneme activation timing patterns described above were due to word frequency effects and not to some other unanticipated factor, all 64 simulations were also run using $k_i = 1.0$ for all words (see section 3 for a description of k_i in equation (3.6)). In this situation word frequencies are completely ignored by the model and can have no impact on simulation results. The $t_{.5}$ results of these simulations are presented in the rightmost two columns of Table 4. In this situation, activation of phoneme nodes of high-frequency irregular words ($t_{.5} = 38.9$) was no longer similar to that of high-frequency regular words ($t_{.5} = 32.5$) but was much more similar to that of low-frequency irregular words ($t_{.5} = 40.3$). In contrast to the previous simulations which included word frequencies, analysis of variance of average phoneme $t_{.5}$ values now indicated that only the main effect of regularity (F 3,60 $= 84.06$, $p < 0.0001$) was highly significant whereas the main effects of frequency (F 3,60 $= 0.12$, $p < .73$) and interaction of regularity and frequency (F 3,60 $= 2.28$, $p < .14$) were not significant. Based on the Student Newman Keuls test, when k_i was constant ($k_i = 1.0$) a significant difference in $t_{.5}$ values of phonemes was found for regular words versus irregular words ($p < 0.005$) while no significant difference was observed in $t_{.5}$ values of phonemes for high-frequency versus low-frequency words. As anticipated, the $t_{.5}$ values for word activation were approximately the same among all four word groups (rightmost column, table 4).

The performance of the print-to-sound model was also tested using non-words as input. The set of non-words consisted of arbitrarily selected, two-syllable words which were not contained in the model's lexicon. In about half the cases of attempted runs with non-words, one and sometimes two or three phonemes never attained full activation because equilibrium was achieved without resolution of a single winner among two or three competing phonemes. For those ten non-words for which winner-takes-all behavior was attained for all phonemes, the mean $t_{.5}$ value for phonemes was 66.6 when k_i was based on word frequency and 68.6 when k_i was constant (see bottom row, Table 4). In both instances, these were signficantly greater $t_{.5}$ values than for the four sets of test words.

A major factor which appeared to influence $t_{.5}$ values for phonemes of non-words was the graphemic/phonemic profile of the non-word's orthographic neighbors contained in the 1030-word lexicon. For example, using graphemes of the non-word "lament" as input (considered a "non-word" since it is not contained in the model's lexicon) produced significant activation of such orthographic neighbors as "latter," "moment," "patent," and "talent." Because these activated words have several conflicting phonemes, more time is required for a winning phoneme to emerge from the competition.

Often for vowel graphemes (which generally had lower connection weights to phonemes than consonant graphemes), orthographic neighbors of non-words influenced the selection of the winning phonemes more strongly than did the probabilities of grapheme-to-phoneme correspondences. For example, using the graphemes of the 'non-word' "cargo" as input resulted in significant activation of the orthographic neighbor "carbon." As a result the winning phoneme for the second grapheme A was /ah/, the second phoneme in "carbon," even though the probabilities (connection weights) for realization of grapheme A as phonemes /ae/ (probability .54), /uh-/ (.19) and /ay/ (.13) all significantly exceeded the probability of /ah/ (.08).

4.4 Simulations with lesioned model

As noted in section 2, part of the support for the dual-route hypothesis of reading aloud comes from studying patients with various forms of acquired dyslexia. The print-to-sound connectionist model described in this paper cannot be related directly to some of the data from these studies because of its circumscribed nature (no semantic influences, no incorrect segmentation of letters into graphemes, etc.). However, it is possible to examine the behavior of a "damaged" model where only one of its two routes is functioning usefully. After completing the simulations described in the previous section, a number of simulations were undertaken where either the lexical or the GPC route was rendered nonfunctional.

Some aspects of "phonological dyslexia" were simulated by disabling the GPC route of the model while leaving the lexical route intact. This was implemented by always setting in_{iG} in equations (3.7) and (3.8) to a constant value of 0.5 for the GPC route for all phoneme nodes in each of the phoneme sets. That is, input from grapheme nodes in G_1 (first instance of a grapheme set) to each of the phoneme nodes in P_1 (first instance of a phoneme set) was clamped at a constant 0.5; the input from grapheme nodes in G_2 to each of the 48 phoneme nodes in P_2 was clamped at 0.5, etc. This universal 0.5 input to phonemes along the GPC route could be interpreted as meaning "no information," so the lexical route entirely determined phoneme activation. In other words, for a given phoneme set, the lexical route inputs forming in_{iW} were the sole source of determining which phoneme became the "winner."

The graphemes of the original set of 64 words were again used as input and simulations were run until a phoneme node in all positions attained 0.999 activation. For all 64 words the correct set of phoneme nodes became fully activated in the same winner-takes-all manner as was observed with the original intact model. The mean $t_{.5}$ value for the 16 high-frequency regular words was 38.2 iterations which was very close to the mean $t_{.5}$ value of 38.3 iterations for the 16 high-frequency irregular words. Likewise, the mean $t_{.5}$ value for the 16 low-frequency regular words (47.3 iterations) was very similar to the mean $t_{.5}$ value for 13 low-frequency irregular words (50.6 iterations). Thus, as might be anticipated with only the lexical route intact, only word frequency and not regularity affected the time for phoneme activation in these simulations. This result is consistent with reported behavior of phonological dyslexic patients, who are unaffected by regularity but read high-frequency better than low-frequency words (e.g., [22]).

Five of the original ten non-words which previously had winning phonemes attaining an activation level of 1.0 with the intact model were arbitrarily selected. Their graphemes were also used as input to the lesioned model. In this case, only partial activation of phonemes occurred when the network reached equilibrium (less than 10^{-3} change in any phoneme activation over a period of 50 iterations). The phonemes which were partially activated never exceeded 0.30 activation and most were less than 0.20. These phonemes corresponded to phonemes of words in the model's lexicon which were orthographic neighbors of the input non-word. Often there were several such words which became partially activated, hence two and sometimes three phonemes in a given phoneme set would remain partially activated at equilibrium, with no clear winner-takes-all phenomenon. For example, using the graphemes of the non-word "compile" as input (since "compile" is not contained in the 1030-word lexicon of the model it is considered to be a "non-word") produced partial activation of several of its orthographic neighbors such as "combine," "compel," and "hostile." Thus, when the network reached equilibrium, partial activation was found for both /p/ and /b/ in the fourth phoneme position, both /eh/ and /ai/ in the fifth phoneme position and both /l/ and /n/ in the sixth phoneme position.

Some aspects of "surface dyslexia" were simulated by disabling the lexical route of our original model while leaving the GPC route intact. This was implemented by clamping in_{iw} to a constant 0.5 in equations (3.7) and (3.8) as input to all phonemes nodes from word nodes. In this situation the grapheme-to-phoneme route inputs in_{iG} were the sole source of determining which phoneme became the "winner" in each phoneme set.

Graphemes for the entire original set of 64 test words were again used as input to the model. For all of the regular words, the correct set of phonemes became totally activated in the usual winner-takes-all manner. The mean $t_{.5}$ value for phonemes of the 32 regular words was 27.1. Graphemes for irregular words also produced winner-takes-all behavior, but did not activate the correct set of phonemes. Rather, the fully activated phoneme nodes were those which had the highest probability among all of the possible grapheme-to-phoneme correspondences. For example, using the five graphemes D E S I G N of the word "design" as input resulted in the winning phonemes being /d eh s ih n/, the set of phonemes with the highest conditional probabilities, instead of the correct /d ih z ai n/. Table 5 lists nine examples of irregular words and the "regularized" winning phonemes (versus the correct phonemes) when run with this model. The mean phoneme $t_{.5}$ value for the 32 irregular words was 28.4. These "regularization" errors are precisely the type of error produced when surface dyslexics attempt to read irregular words (e.g., [7]).

Graphemes for the same five non-words used with the "phonological dyslexia" model were again used as input. These simulations also produced winner-takes-all behavior with a "regularized" pronunciation. Just as with irregular words, fully activated phoneme nodes were always those which had the highest probability among all GPC frequencies. The mean phoneme $t_{.5}$ value was 27.9.

5. Discussion

The print-to-sound connectionist model described in this paper is based on an indirectly interactive dual-route associative network derived from data published in the psychological literature. A competitive activation mechanism obviates the need for inhibitory connections between nodes representing mutually-exclusive outcomes (word and phoneme nodes). The results of studying this model are of interest both because of their implications for implementing competitive dynamics in connectionist models and because of the support they provide for an interactive dual-route hypothesis.

Testing of the intact model with graphemic input corresponding to 64 words of varying frequency and regularity always resulted in clear-cut winner-takes-all behavior by correct nodes in every case. Testing the model with input graphemes forming "non-words" sometimes resulted in one or two phoneme positions failing to establish a clear "winner" at equilibrium. This occurred when the partially activated word nodes in the model's limited lexicon conflicted concerning the correct phoneme and the winner was "too close to call." In other studies with competitive activation mechanisms, it has proven possible to sharpen competitive effects and force selection of a winner in close outcomes by suitable alterations of the activation rule [41,38]. It is likely that such an approach could be used to improve the print-to-sound model's performance with non-word graphemic input, but this possibility was not explored in the current study.

These results provide encouragement concerning the direct applicability of competitive activation methods to the large, complex associative networks

Word	APPROVE					
Graphemes	A	PP	R	O-E	V	
Winning Phonemes	ae	p	r	o	v	
Correct Phonemes	uh-	p	r	oo	v	
Word	BEHIND					
Graphemes	B	E	H	I	N	D
Winning Phonemes	b	eh	h	ih	n	d
Correct Phonemes	b	ih	h	ai	n	d
Word	DECLARE					
Graphemes	D	E	C	L	A-E	R
Winning Phonemes	d	eh	k	l	ay	r
Correct Phonemes	d	ih	k	l	eh	r
Word	DESIGN					
Graphemes	D	E	S	I	GN	
Winning Phonemes	d	eh	s	ih	n	
Correct Phonemes	d	ih	z	ai	n	
Word	MAJOR					
Graphemes	M	A	J	O	R	
Winning Phonemes	m	ae	dj	o	r	
Correct Phonemes	m	ay	dj	er	r	
Word	MONKEY					
Graphemes	M	O	N	K	EY	
Winning Phonemes	m	o	n	k	ee	
Correct Phonemes	m	uh+	ng	k	ee	
Word	REMIND					
Graphemes	R	E	M	I	N	D
Winning Phonemes	r	eh	m	ih	n	d
Correct Phonemes	r	ih	m	ai	n	d
Word	TREASURE					
Graphemes	T	R	EA	S	U-E	R
Winning Phonemes	t	r	ee	s	yu	r
Correct Phonemes	t	r	eh	zh	er	r
Word	UNION					
Graphemes	U	N	I	O	N	
Winning Phonemes	uh+	n	ih	o	n	
Correct Phonemes	yu	n	y	uh-	n	

Table 5: Phoneme activation patterns for irregular words when only the GPC route is intact.

of interest in cognitive science and AI. Previous connectionist models have avoided diffuse network saturation as activation spreads throughout a network in a variety of ways. In some cases, decay has been used [18]. More commonly, lateral inhibitory links have been added to an otherwise excitatory network [14,20,34], perhaps reflecting the rather discouraging observation that "lateral inhibition seems to have fewer disadvantages" than alternative mechanisms [57, p. 55]. Competitive activation mechanisms offer a third alternative with some significant advantages: lateral inhibitory links are unnecessary (over one million avoided in the print-to-sound network), multiple fully-active "winners" can be sustained when appropriate, etc. (see section 2). Of course, much work remains to be done to extend this approach. Important issues needing study in future research include how to combine competitive activation methods with symbol-processing approaches [29], gating of node output on semantically labeled links (e.g., inheritance in semantic category hierarchies), and methods for systematically deriving competitive activation rules.

The print-to-sound transformation model is also of interest because of its ability to replicate, at least qualitatively, a number of previously-observed behavioral phenomena. This is particularly striking in the context of the restricted nature of the implementation (only 1030 words in the lexicon, no segmentation analysis, etc.). When words are treated as if they all have the same frequency, the model "pronounces" (generates phonemic representation for) regular words more quickly than irregular words. In contrast, when word frequency is factored in as an influence on word activation, only those irregular words having a low frequency are found to be "pronounced" slower than regular words; high-frequency irregular words are pronounced at the same rate as regular words. All of these results are consistent with observations made with normal readers [49,28].

Only a few other connectionist models of print-to-sound transformation have been reported to date and most of these are quite different in their goals and methods. For example, NETtalk is fundamentally different in that it uses a distributed representation of information, does not use a competitive activation mechanism, does not explicitly represent a dual-route network, and applies error backpropagation to learn connection weights [50]. The most similar previous work known to the authors is Brown's word naming model [6]. It is difficult to compare Brown's model with the one described here because very little information on the actual implementation of his model is presented. However, Brown's model differs from ours in that it uses explicit inhibitory connections, and rather than using grapheme-to-phoneme correspondences, it maps multiple input segments (single, double, triple letters) into various phonological codes. In addition, the lexicon it uses is limited to an unspecified number of four-letter words with frequencies of either "high" or "low."

There are currently a number of hypotheses concerning the cognitive mechanisms involved in reading aloud. Many of these involve what appear to be complex interactions, such as the relationship between frequency and regularity in determining reading times, as discussed above. The nature of these interactions are difficult to determine simply from experimentation with research subjects. One of the values of detailed computational models of cognitive processes is that they both force one to be explicit about an implementation and permit one to determine whether or not anticipated behaviors truly can arise from a given manipulation. It is precisely in this sense that the correspondence of the behavior of the print-to-sound model to a number of phenomena observed in both normal and dyslexic readers provides support for the general consistency of the dual-route hypothesis.

Acknowledgements

This work was supported in part by ONR Award N00014-85-K0390 and in part by NSF Award IRI-8451430 with matching funds from AT&T Information Systems. Dr. Berndt holds an NIH RCDA Award K04-NS00851. Technical assistance in preparation of this manuscript from D. Calo is gratefully acknowledged.

Appendix

High-frequency Regular	High-frequency Irregular	Low-frequency Regular	Low-frequency Irregular
Indeed	Design	Daylight	Accord
Meeting	Union	Frighten	Approve
Simple	Future	Plainly	Odor
Maybe	Above	Feeble	Onion
Feeling	Unit	Needle	Idle
Little	Foreign	Ample	Array
Middle	Report	Lazy	Torture
Highly	Private	Apple	Absorb
Training	Open	Swiftly	Armor
District	Behind	Railway	Treasure
Clearly	Market	Dismay	Depart
Standing	Labor	Lately	Resort
Inside	Major	Kitten	Monkey
Likely	Color	Upright	Vapor
Nearly	Forward	Whistle	Declare
Figure	Record	Invade	Remind

References

[1] D. A. Allport, "On Knowing the Meaning of Words We are Unable to Report: The Effects of Visual Masking" in *Attention and Performance*, **6**, ed. S. Dornic (Lawrence Erlbaum Associates, Hillsdale, NJ, 1987)

[2] S. Amari, and M. Arbib, eds. *Competition and Cooperation in Neural Nets* (Springer-Verlag, New York, 1982).

[3] M. Beauvois and J. Derousne, "Phonological Alexia: Three Dissociations," *Journal of Neurology, Neurosurgery, and Psychology*, **42** (1979) 1115–1124.

[4] R. Berndt, J. Reggia, and C. Mitchum, "Empirically-Derived Probabilities for grapheme-to-phoneme Correspondences in English," *Behavior Research Methods, Instruments and Computers*, **19** (1987) 1–9.

[5] G. Blelloch and C. Rosenberg, "Network Learning on the Connection Machine," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, ed. J. McDermott, (Morgan Kaufmann Publishers, Los Altos, CA, 1987).

[6] G. Brown, "Resolving Inconsistency: A Computational Model of Word Naming," *Journal of Memory and Language*, **26** (1987) 1–23.

[7] D. Bub, A. Canceliere, and A. Kertesz, "Whole-word and Analytic Translation of Spelling to Sound in a Non-Semantic Reader" in *Surface Dyslexia: Cognitive and Neuropsychological Studies of Phonological Reading*, eds. K. E. Patterson, J. C. Marshall and M. Coltheart (Lawrence Erlbaum Associates, London, 1985).

[8] M. Coltheart, "Lexical Access in Simple Reading Tasks" in *Strategies of Information Processing*, ed. G. Underwood (Academic Press, London and New York, 1978).

[9] M. Coltheart, "Reading, Phonological Recoding and Deep Dyslexia" in *Deep Dyslexia*, eds. M. Coltheart, K. Patterson, and J. C. Marshall (Routledge and Kegan Paul, London, 1980).

[10] M. Coltheart, "Writing Systems and Reading Disorders" in *Orthographies and Reading*, ed. L. Henderson (Lawrence Erlbaum Associates, London, 1984).

[11] M. Coltheart, "Cognitive Neuropsychology and the Study of Reading" in *Attention and Performance XI*, eds. M. I. Posner and O. S. M. Marin (Lawrence Erlbaum Associates, Hillsdale, NJ, 1985) 3–37).

[12] M. Coltheart, "The Psychology of Reading," *Attention and Performance, 12* (Lawrence Erlbaum Associates, London, 1987).

[13] M. Coltheart, J. Masterson, and S. Byng, "Surface Dyslexia," *Quarterly Journal of Experimental Psychology*, **35** (1983) 469–495.

[14] G. Cottrell and S. Small, "A Connectionist Scheme for Modelling Word Sense Disambiguation," *Cognition and Brain Theory*, **6** (1983) 89–120.

[15] Commentary on [30]: *Behavioral and Brain Sciences*, **8** (1985) 705–739 and **10** (1987) 331–336.

[16] C. L. D'Autrechy and J. Reggia, "The Mirrors/II Simulator," *Proceedings of the Twentieth Annual Simulation Symposium*, (1987) 121–132.

[17] C. L. D'Autrechy, J. Reggia, G. Sutton, and S. Goodall, "A General Purpose Simulation Environment for Developing Connectionist Models," *Simulation*, **51** (1988) 5–19.

[18] G. Dell, "A Spreading Activation Theory of Retrieval in Sentence Production," *Psychological Review*, **93** (1986) 283–321.

[19] A. Ellis, *Reading, Writing and Dyslexia: A Cognitive Analysis* (Lawrence Erlbaum Associates, Hillsdale, NJ, 1984).

[20] J. Feldman and D. Ballard, "Connectionist Models and Their Properties," *Cognitive Science*, **6** (1982) 205–254.

[21] M. Friendly, P. E. Franklin, D. Hoffman, and D. C. Rubin, "The Toronto Word Pool: Norms for Imagery, Concreteness, Orthographic Variables, and Grammatical Usage for 1,080 Words," *Behavior Research Methods and Instrumentation*, **14** (1982) 375–399.

[22] E. Funnel, "Phonological Processes in Reading: New Evidence from Acquired Dyslexia," *British Journal of Psychology*, **74** (1983) 159–180.

[23] R. J. Glushko, "The Organization and Activation of Orthographic Knowledge in Reading Aloud," *Journal of Experimental Psychology: Human Perception and Performance*, **6** (1979) 674–691.

[24] P. B. Gough, and M. J. Cosky, "One Second of Reading Again" in *Cognitive Theory II*, eds. N. J. Castellan, D. B. Pisoni, and G. H. Potts (Lawrence Erlbaum Associates, Hillsdale, NJ, 1977).

[25] S. Grossberg, *Studies of Mind and Brain* (Reidel Publishing Co., Norwell, MA, 1982).

[26] P. R. Hanna, J. S. Hanna, R. E. Hodges, and E. H. Rudorf, *Phoneme-Grapheme Correspondences as Cues to Spelling Improvement* (U.S. Department of Health, Education, and Welfare, Washington, D.C., 1966).

[27] L. Henderson, *Orthography and Word Recognition in Reading* (Academic Press, London, 1982).

[28] L. Henderson, "Issues in the Modelling of Pronunciation Assembly in Normal Reading" in *Surface Dyslexia: Neuropsychological and Cognitive Studies of Phonological Reading*, eds. K. E. Patterson, J. C. Marshall, and M. Coltheart (Lawrence Erlbaum Associates, Inc., London, 1985).

[29] J. Hendler, "Marker Passing and Microfeatures" in *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, ed. J. McDermott (Morgan Kaufmann, Los Altos, 1987) 151–154.

[30] G. Humphreys, and L. Evett, "Are There Independent Lexical and Nonlexical Routes in Word Processing? An Evaluation of the Dual-Route Theory of Reading," *Behavioral and Brain Sciences*, **8** (1985) 689–705.

[31] J. Kay,and D. Bishop, "Anatomical Differences between Nose, Palm, and Foot, or the Body in Question: Further Dissection of the Processes of Sub-Lexical Spelling-Sound Translation" in *Attention and Performance*, ed. M. Coltheart (Ed.) (Lawrence Erlbaum Associates, London, 1987).

[32] J. Kay and A. Marcel, "One Process, Not Two, in Reading Aloud: Lexical Analogies Do the Work of Nonlexical Rules," *Quarterly Journal of Experimental Psychology*, **33A** (1981) 397–413.

[33] H. Kucera and N. Francis, *Computational Analysis of Present-Day English* (Brown University Press, Providence, RI, 1967).

[34] J. McClelland and D. Rumelhart, "An Interactive Activation of Model of Context Effects in Letter Perception. Part I: An Account of Basic Findings," *Psychological Review*, **88** (1981) 375–407.

[35] J. Morton and K. Patterson, "A New Attempt at an Interpretation, or, an Attempt at a New Interpretation" in *Deep Dyslexia*, eds. M. Coltheart, K. Patterson, and J. C. Marshall (Routledge and Kegan Paul, London, 1980).

[36] F. Newcombe and J. Marshall, "Varieties of Acquired Dyslexia: A Linguistic Approach" in *Seminars in Neurology*, **4** (1984) 181–195.

[37] K. Patterson and V. Coltheart, "Phonological Processes in Reading: A Tutorial Review" in *Attention and Performance*, ed. M. Coltheart (Lawrence Erlbaum Associates, London, 1987).

[38] Y. Peng and J. Reggia, "A Probabilistic Causal Model for Diagnostic Problem Solving," *IEEE Transactions of Systems, Man, and Cybernetics*, **17** (1987) 146–162, 395–406.

[39] Y. Peng, and J. Reggia, *A Connectionist Model for Diagnostic Problem Solving, IEEE Trans. on Systems, Man and Cybernetics*, 1989, in press.

[40] J. Reggia, "Virtual Lateral Inhibition in Parallel Activation Models of Associative Memory," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, **1** (1985) 244–248.

[41] J. Reggia, "Properties of a Competition-Based Activation Mechanism in Neuromimetic Network Models," *Proceedings of the First International Conference on Neural Networks*, Vol. II, San Diego, CA (1987) 131–138

[42] J. Reggia, D. Nau, and P. Wang, "Diagnostic Expert Systems Based on a Set Covering Model," *International Journal of Man-Machine Studies*, **19** (1983) 437–460.

[43] J. Reggia, D. Nau, and P. Wang, "A Formal Model of Diagnostic Inference," *Information Sciences*, **37** (1985) 227–285.

[44] J. Reggia, and G. Sutton, "Self-Processing Networks and Their Biomedical Implications," *Proc. IEEE*, **76** (1988) 680–692.

[45] M. B. Rosson, "The Interaction of Pronunciation Rules and Lexical Representations in Reading Aloud," *Memory and Cognition*, **13(1)** (1985) 90–99.

[46] D. Rumelhart and J. McClelland, "An Interactive Activation Model of Context Effects in Letter Perception. Part 2: The Contextual Enhancement Effect and Some Tests and Extensions of the Model," *Psychological Review*, **89** (1982) 60–94.

[47] D. Rumelhart, J. McClelland, and the PDP Research Group, *Parallel Distributed Processing - Volume 1: Foundations* (MIT Press, Cambridge, MA, 1986).

[48] E. Saffran, "Acquired Dyslexia: Implications for Models of Reading" in *Reading Research: Advances in Theory and Practice*, **IV**, eds. G. MacKinnon and T. Waller (Academic Press, 1984).

[49] M. S. Seidenberg, G. S. Waters, M. A. Barnes, and M. K. Tanenhaus, "When Does Irregular Spelling or Pronunciation Influence Word Recognition?," *Journal of Verbal Learning and Verbal Behavior*, **23** (1984) 383–404.

[50] T. Sejnowski and C. Rosenberg, "Parallel Networks that Learn to Pronounce English Text," *Complex Systems*, **1** (1987) 145–168.

[51] S. Seidman, P. Wang, and J. Reggia, "Analysis of a Competition-Based Connectionist Model," *Proceedings of the First International Conference on Neural Networks*, Vol. II, San Diego, CA (1987) 139–145.

[52] P. Smolensky, "On the Proper Treatment of Connectionism," *Behavioral and Brain Sciences*, in press.

[53] K. E. Stanovich and D. W. Bauer, "Experiments on the Spelling-to-Sound Regularity Effect in Word Recognition," *Memory and Cognition*, **6** (1978) 410–415.

[54] M. Tagamets, and J. Reggia, "A Data Flow Implementation of a Competition-Based Connectionist Model," *Journal of Parallel and Distributed Computing*, 1989, in press.

[55] R. L. Venezky, *The Structure of English Orthography* (Mouton, The Hague, 1970).

[56] *Webster's New Collegiate Dictionary* (8th ed.) (Merriam-Webster, Springfield, MA, 1981).

[57] D. Waltz and J. Pollack, "Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation," *Cognitive Science*, **9** (1985) 51–74.