Cognitive Neuropsychology
2001, Vol. 18, No. 1, 71-92

Are There Orthographic Impairments In Phonological Dyslexia?

Michael W. Harm
Center for the Neural Basis of Cognition
Carnegie Mellon University

Mark S. Seidenberg
Neuroscience Program
University of Southern California

Two hypotheses have been advanced concerning the basis of acquired phonological dyslexia. According to the dual-route model, the pattern derives from impaired grapheme-phoneme conversion. According to the phonological impairment hypothesis, it derived from impaired representation and use of phonology. Effects of graphemic complexity and visual similarity observed in studies by Howard and Best (1996), orthographic effects on phoneme counting (Berndt, Haendiges, Mitchum, & Wayland, 1996) and data from patient LB (Derouesné & Beauvois, 1985) have been taken as evidence for an orthographic impairment in phonological dyslexia and therefore against the impaired phonology hypothesis (Coltheart, 1996). We present a computational simulation, results of two behavioral studies and a critical analysis of the MJ and LB data which suggest that the "orthographic" deficits in such patients arise from phonological impairments that interact with orthographic properties of stimuli.

Introduction

Acquired Phonological Dyslexia, a pattern of impaired reading that is observed following some types of neuropathology, (Beauvois & Derouesné, 1979; Derouesné & Beauvois, 1979), has played an important role in theories of visual word recognition. The pattern is characterized by a primary impairment in reading nonwords such as NUST and stands as a counterpart to surface dyslexia (Patterson, Marshall, & Coltheart, 1985) in which the patient shows a primary impairment in reading words with irregular spelling to sound correspondences (exception words such as HAVE).

These complementary reading impairments have been taken as evidence for the "dual-route" theory developed by Coltheart and colleagues (e.g., Coltheart, Davelaar, Jonasson, & Besner, 1977; Coltheart, Curtis, Atkins, & Haller, 1993), in which there are two independent pathways from print to sound: one involving word-specific knowledge, and the other involving the application of grapheme to phoneme conversion (GPC) rules (see Figure 1). Exception words are assumed to be read via the word-specific system; surface dyslexia is then interpreted as an impairment in this process. Nonwords are read by means of GPC rules; phonological dyslexia is attributed to an impairment in this process. Thus the model provides an elegant account of the main features of the two complementary impairments.

Our focus in this article is on an alternative account of phonological dyslexia inspired by the "triangle" connectionist model of reading (Seidenberg & McClelland, 1989; Plaut, McClelland, Seidenberg, & Patterson, 1996; Harm, 1998) and related work (see Figure 2). An account of surface dyslexia within this framework was presented in Plaut et al. (1996) and Seidenberg (1995). Several considerations suggest that phonological dyslexia is caused by an impairment in the representation of phonological information rather than grapheme-phoneme conversion. We shall call this the *phonological impairment hypothesis* (Patterson & Marcel, 1992; Farah, Stowe, & Levinson, 1996; Patterson, Suzuki, & Wydell, 1996; Sasanuma, Ito, Patterson, & Ito, 1996; Manis, Seidenberg, Doi, McBride-Chang, & Peterson, 1996; Seidenberg, 1995; Harm & Seidenberg, 1999). According to this hypothesis, the advantage of word reading over nonword reading derives from nonwords having a

---

Michael W. Harm, Center for the Neural Basis of Cognition, Carnegie Mellon University; Mark S. Seidenberg, Neuroscience Program and Department of Psychology, University of Southern California.

This work was supported by National Institute of Mental Health (NIMH) grant 47566, National Institute of Child Health and Human Development (NICHD) grant 29891, a Research Scientist Development Award (NIMH 01188) and a National Research Service Award from the National Institute of Deafness and Other Communication Disorders (NIDCD), National Institutes of Health, grant F32-DC00425.

We thank David Plaut, Robert Thornton and Jason Zevin for helpful comments on an earlier draft of this paper.

Correspondence concerning this article should be addressed to Michael W. Harm, Center for the Neural Basis of Cognition, 115 Mellon Institute, 4400 Fifth Avenue, Pittsburgh, Pennsylvania, 15213. Electronic mail may be sent to mharm@cnbc.cmu.edu.

less stable phonological representation than words; therefore phonological impairment yields more errors on nonword reading than word reading. There are several complementary reasons why nonwords would produce a less stable phonological representation: the phonological forms of nonwords are (by definition) less familiar to the reading system, they are unable to receive complementary activation from the semantic system because they do not activate semantics as strongly as words do, and finally, words (unlike nonwords) can in principle activate semantics directly in the triangle model without phonological mediation (via the orth→sem connections in Figure 2), which can in turn provide additional support to phonology via the sem→phonology connections. All of these reasons potentially give words stronger activation of phonological codes than nonwords.

One source for this hypothesis was the observation by Besner, Twilley, McCann, and Seergobin (1990) that Seidenberg and McClelland's (1989) connectionist model of word recognition behaved like a phonological dyslexic: it computed the correct pronunciations for words but did relatively poorly on nonwords. Seidenberg and McClelland (1990) and Plaut et al. (1996) provided analyses showing that the source of those nonword errors was limitations of the phonological representation used in the Seidenberg and McClelland (1989) model. Plaut et al. (1996) presented simulations with an improved phonological representation that yielded better nonword performance. Thus the models demonstrate how impaired generalization can arise from a phonological impairment rather than impaired use of GPC rules.

A second observation consistent with this hypothesis is that an overwhelming majority of patients with acquired phonological dyslexia have also exhibited impaired use of phonological information on tasks unrelated to nonword reading. In the 1996 issue of Cognitive Neuropsychology devoted to phonological dyslexia, a total of 18 patients were reported. As noted by Coltheart (1996), all of them were phonologically impaired. Of the other phonological dyslexics reported in the literature whose phonological abilities were tested, all but one were impaired. The dual route model must treat this broader phonological impairment as one that happens to co-occur with the primary deficit in grapheme-phoneme conversion. In contrast, this co-occurrence follows directly from the phonological impairment hypothesis: the patients have a phonological impairment that affects both nonword reading and performance on other tasks involving this information.

Additional evidence supporting the phonological impairment hypothesis is provided by studies of developmental dyslexia. Developmental phonological dyslexia is a variant of the classic acquired case, in which children learning to read show an impairment in nonword reading (e.g., Castles & Coltheart, 1993; Manis et al., 1996). Many studies have found that children exhibiting impaired nonword reading also exhibit significantly worse performance on other phonological tasks (Manis et al., 1996; Stanovich, Siegel, & Gottardo, 1997). This analysis of developmental phonological dyslexia is consistent with the extensive literature concerning the linkage between phonological ability and reading development (see Adams, 1990, for a review). Phonological analysis is essential for the formation of componential representations that can support generalizations of the spelling to sound regularities: it is only by being able to hear the sound overlap of RAT/RATE and BIT/BITE that one can abstract the generalization of the effect of the final E on the vowel[1].

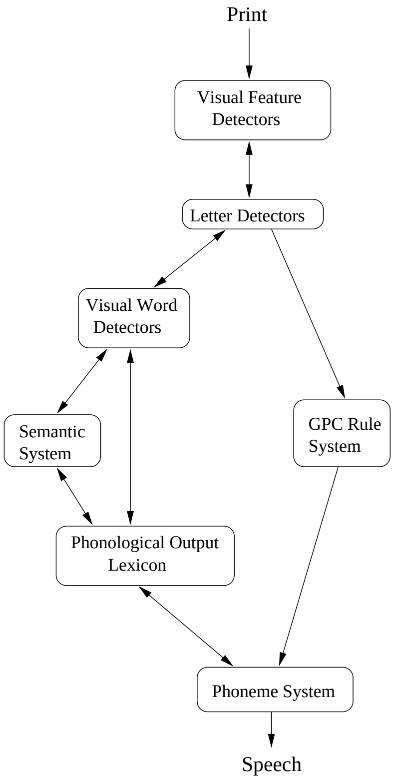

Figure 1. The DRC model of word recognition (Coltheart et al., 1993).

Recent computational models of normal and impaired reading provide additional evidence concerning the causal relation between poor phonological representations and reading impairments. Harm and Seidenberg (1999) simulated developmental phonological dyslexia by impairing the

---

[1] A very similar observation holds for the development of inflectional morphology. In a very similar vein, the vast majority of children who exhibit developmental problems learning morphological regularities also exhibit phonological processing impairments (Bishop, 1992). Joanisse and Seidenberg (1999) present simulations relevant to impairments in inflectional morphology; they find that impairing the phonological component of a model yields impaired ability to generalize the past tense inflection to novel forms.

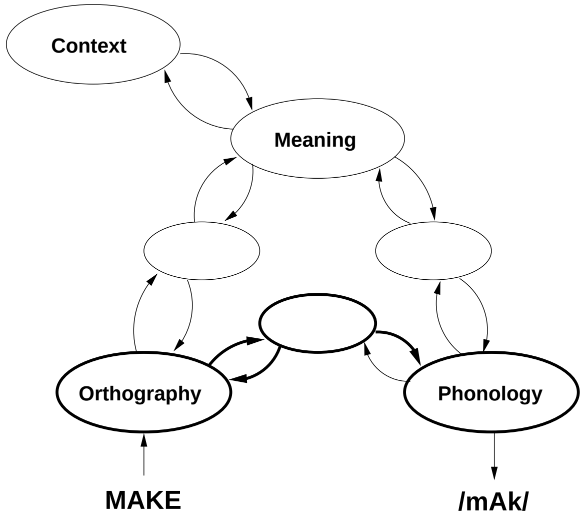

Figure 2. The Seidenberg and McClelland (1989) model of word recognition. Implemented pathways are shown in bold.

development of phonological representations in different degrees. Mild phonological damage produced a selective impairment in nonword reading, with normal word reading. More severe phonological damage produced a mixed case, in which nonword and exception word reading were both below normal levels, with nonword reading showing the more severe impairment.

Evidence Against the Phonological Impairment Hypothesis

Coltheart (1996) argued that two types of evidence contradict the phonological impairment hypothesis. One is the deficits in orthographic processing observed in some cases (Derouesné & Beauvois, 1985; Howard & Best, 1996). These include effects of graphemic complexity and orthographic similarity. The other is data from patient LB (Derouesné & Beauvois, 1985), who was said to exhibit impaired nonword reading without an accompanying phonological deficit. We consider these in turn.

Coltheart (1996) argues that sensitivity to graphemic complexity is evidence against a purely phonological cause of reading impairments (p. 757):

> One example of sensitivity to such a variable has already been mentioned: Two of the patients studied by Derouesnè and Beauvois (1979) were worse at reading nonwords that contained two-letter graphemes than nonwords that did not. This difference could not be a phonological effect; it must be orthographic, and that remains true even if these two patients had demonstrable phonological impairments. LB (Derouesnè and Beauvois, 1985) and the developmental case described by Howard and Best (this issue) also showed this influence of graphemic complexity on nonword reading accuracy.

Similar logic is used by Berndt et al. (1996) in analyzing a graphemic segmentation task. Subjects were asked to say how many phonemes are in a string, such as AUK, which has two phonemes, or VAD, which has three. Berndt et al. (1996) argue that this provides a test of orthographic/graphemic analysis independent of phonological processing. They found that all of their subjects were impaired in both ability to segment graphemes and phonological processing. They interpreted the data as evidence that phonological dyslexia involves multiple independent impairments.

Derouesné and Beauvois (1985) also argued that the orthographic effects (higher rates of errors on graphemically complex items) seen in patient LB's performance showed an impairment in a graphemic processing stage. Howard and Best (1996) tested a phonological dyslexic, MJ, and discovered an effect of graphemic complexity in her ability to read nonwords (e.g. complex words like CHACK versus simple words like BEM). This effect was argued to show the subject had an impairment in orthographic analysis above and beyond her reported phonological impairment, because phonological factors were matched: both conditions utilized words with CVC phonological structure.

Phonological dyslexic patients also exhibit an effect of visual similarity on the processing of pseudohomophones. In many studies of normals and patients, pseudohomophones such as BRANE were named faster than nonpseudohomophones such as BRONE. The further finding is that in some studies this effect was modulated by the extent to which the pseudohomophone is visually similar to the target homophonous word, e.g., GERL is read more accurately than PHOCKS. Patient LB (Derouesné & Beauvois, 1985) exhibited this effect, as did patient MJ (Howard & Best, 1996). Howard and Best argue that this effect is consistent with the dual route model but not single mechanism models (e.g., Plaut et al., 1996). The account of the pseudohomophone advantage given by Patterson et al. (1996), in which a phonological pattern that maps onto a stable semantic representation receives support from semantics, is said to disallow orthographic effects because there is no orthography in the semantic-to-phonological attractor.

Finally, the existence of a phonological dyslexic patient without a phonological impairment is argued by Coltheart (1996) to provide evidence against the phonological impairment hypothesis. Patient LB, it is claimed, shows a clear nonword reading impairment with normal phonological processing abilities.

Taken together, these observations call into question the hypothesis that phonological dyslexia derives solely from a phonological impairment. Moreover, the evidence from patient LB raises the possibility that the phonological impairments seen in every other case of phonological dyslexia were co-occurring deficits rather than causal.

The evidence that an orthographic processing impairment accounts for the patterns of behavior seen in phonological dyslexics (e.g., graphemic complexity effects, visual similarity modulation of the pseudohomophone effect) all arise from the same kind of manipulation: phonological factors (complexity, or lexical status) are held constant and orthographic factors are manipulated. When an effect is seen, it is argued that it cannot arise from purely phonological impairments because the phonological factors are the same in both conditions.

This line of reasoning follows naturally within the DRC model of word recognition, which posits a stagelike processing of spelling into sound. In this model, orthographic word forms are parsed left to right into a series of graphemes. These graphemes are converted to phonemes by the grapheme to phoneme conversion system. Because graphemes are pre-parsed, the phonological system does not "know" how complex the orthographic pattern was; the phonemes delivered to phonology for FOX are the same whether the input was FOX or PHOCKS. So differences in performance on a task (such as producing the pronunciation, or counting the resultant phonemes, or activating a semantic representation from a wordlike sound pattern) cannot be explained via phonological impairments, as the phonological system in the DRC model receives the same input for both PHOCKS and FOX.

However, there are other models of reading in which such conclusions do not logically follow. In interactive PDP models (e.g., Seidenberg & McClelland, 1989; Harm & Seidenberg, 1999), orthographic representations are not "parsed" into graphemes; the phonological output is computed directly from orthographic input. Hence, in these models, because orthography and phonology are connected (either directly or via hidden units), the phonological system does "know" whether the input was PH or F. Impairments in phonological processing will necessarily impact the more difficult items more so than simpler ones.

An analogy can perhaps illustrate what we mean. Consider a game of golf in which the goal for the next shot is to get the ball from its current position to the green. The difficulty of the shot is affected by the distance from the green (e.g., 30 vs. 60 yards away), but the target remains the same. A skilled golfer might make both shots, although the longer shot is more difficult and requires more skill. Suppose now that the target is degraded in some fashion: say the green is smaller or has more divots. The same golfer from the same starting points may make the easier shot but miss the harder one. Thus properties of the target interact with "input" factors such as the ball's current position and the golfer's skill.

In this analogy, the green corresponds to a target phonological pattern, the current position is the "input pattern," and distance from the green corresponds to orthographic complexity, a factor that affects how hard it is to achieve the input-output mapping. It would be erroneous to argue that the condition of the green could not have affected the difficulty of the shot because the target is identical in the two cases. In fact the target is not identical and the effect of degrading it depends on how hard the shot was in the first place.

The main point of course is that orthography and phonology are connected rather than independent in our models. Thus orthographic manipulations have phonological consequences; when phonology is degraded, the "starting position" (degree of orthographic complexity) matters more.

In summary, this analysis suggests that, rather than arising from impaired orthographic analysis, the "orthographic" effects seen in phonological dyslexics arise from a phonological impairment whose effects depend on properties of the orthographic input. In the remainder of this article we present three lines of evidence bearing on this hypothesis. First, we show that the "orthographic" effects evidenced by patients' errors in reading nonwords are also observed in normal subjects' response latencies. We used items from the Howard and Best (1996) studies in an experiment with undergraduate subjects and compared their latencies to patient MJ's errors. In both cases performance is affected by orthographic properties of the stimuli; phonological damage merely magnifies this effect. Second, we describe a connectionist model that produces the conjectured effects. In its normal unimpaired state the model produces behavior like that of normal subjects; introducing a phonological impairment–but leaving the orthographic input untouched–creates the pattern characteristic of MJ. Finally, we examine the evidence concerning patient LB, the only patient said to have a nonword reading impairment with normal phonology. We argue that it is far from clear from their report that LB was phonologically unimpaired, and suggest that biases in their stimuli can account for his performance.

1. Graphemic Complexity Effects

Howard and Best (1996) demonstrated an effect of graphemic complexity in subject MJ's reading accuracies. They utilized three conditions: a simple/3 condition, in which three letter words map one-to-one onto a three phoneme pronunciation (e.g., BEM), a simple/5 condition in which five letters map one-to-one onto a five phoneme pronunciation (e.g., BLISK), and a complex condition in which a five letter spelling pattern maps onto three phonemes (e.g., CHACK). Simple/3 and complex were the crucial conditions: the conditions were equated in terms of phonological complexity in that the targets were always CVCs. What differed was the complexity of the orthographic code.

On our view, nonwords containing more complex correspondences between orthography and phonology will be more difficult, both for normal subjects and the model. Each will be considered in turn.

Experiment 1: Graphemic Complexity Effects in Normals

This experiment tested for an effect of graphemic complexity in naming RTs in normal college undergraduate students.

Method

Subjects. Twenty four members of the University of Southern California community received course credit or were paid $5 for their participation in the experiment. None had a recorded history of reading impairment. All were native speakers of English. One subject's results were excluded from analysis because his error rate was very high (25% of trials) and his RTs were quite large (mean 1380ms).

Stimuli and Design. The materials were the same items used in the Howard and Best (1996) graphemic complexity study, counterbalanced against items from Experiment 2[2]. Subjects were presented with one word at a time on a Macintosh computer screen. Time to name the item was recorded using a Psyscope voice activated relay. An experimenter recorded whether subject named the item correctly or not using a Psyscope button box.

Results

Trials with an RT $< 300ms$ or $> 1200ms$ were coded as equipment failures and were excluded from RT analysis. Trials with a reaction time greater than 2.5 standard deviations greater than or less than the mean for their cell were withheld from analysis. This removed 3.2% of the trials.

The accuracy and reaction time results are presented in Table 1, along with subject MJ's latencies and accuracies. The effect of complexity on reaction time was significant by subjects ($F1(2,44) = 85.5$, $p < 0.001$) and by items ($F2(2,128) = 49.7$, $p < 0.001$). The results pattern in the same way as patient MJ's errors.

It is possible that the reaction time differences obtained in Experiment 1 do not reflect the relative difficulty of the stimuli, but more narrowly reflect the operation of a graphemic parser. This would be consistent with the conclusions of Howard and Best (1996). To test the possibility that a purely phonological impairment could give rise to the patterns of impairments seen in patient MJ, we also used a connectionist model of word recognition based on the original "triangle" formulation from Seidenberg and McClelland (1989) to test for effects of graphemic complexity, and the impact of phonological damage on any such effect.

Simulation 1: MJ Graphemic Complexity

Method

Materials. The model was trained using a set of 6,103 monosyllabic words culled from the Wall Street Journal corpus (Marcus, Santorini, & Marcinkiewicz, 1993). The model contained adjectives, closed class words, adverbs, singular and plural nouns, and present tense, past tense and third person singular verbs.

The phonological representation was a slot based, CCCVCCC structure with 25 binary phonetic features per slot to code each phoneme, yielding 175 phonological features. These features were derived from Chomsky and Halle (1968), with minor modifications.

The semantic representations for the words in the model were derived from the WordNet online semantic database. These features encoded high level semantic information (such as <object> or <living thing>), as well as more low level information (such as <has-a-carburator>). There were no synonyms, although many items had a large number of shared features. Morphological relationships were coded with features such as <plural>; for example, the only difference between the semantic representation of CAT and CATS was the <plural> feature. All features were binary: if a feature was present for a given item, its value was 1.0, otherwise it was 0. A total of 1989 semantic features were used. The number of active features for a word ranged from 1 to 37, with a mean of 7.6 and standard deviation of 4.4. See Harm (1998) for more details.

The orthographic representations were coded using a slot based representation, using binary localist units for letters, fitting into a vowel centered CCCVVCCCCC template. A total of 133 orthographic features were used.

The test items from the graphemic complexity manipulation performed by Howard and Best (1996) were tested against the model.

Architecture. Figure 3 depicts the model architecture. The semantic representations have a set of 50 "cleanup" units which constrain their activity; the phonological representations also have 50 cleanup units. Two sets of 500 hidden units mediate activity between semantics and phonology, one set for each direction. A set of 100 hidden units map between orthography and phonology; the orthographic units are also connected directly to the phonological units. A larger set of 500 hidden units map between orthography and semantics, the orthographic units are also connected directly to the semantic units. These numbers were chosen based on earlier pilot simulations which indicate that they are sufficient to accomplish the task without being overly computationally burdensome.

It should be noted that the model's architecture includes a set of connections mapping orthographic inputs directly onto phonological units, and another set which maps orthographic units onto semantic units. These connections were added to the system mapping orthography to phonology because we have found that such connections can tend to improve generalization performance of a network. Direct connections were added from orthography to semantics chiefly for symmetry; it was less obvious that they would be of any direct benefit to that part of the system (although see the General Discussion where we describe evidence of sublexical semantic reading in the model).

---

[2] We thank David Howard for kindly providing us with the stimuli.

Zorzi, Houghton, and Butterworth (1998) explored a spelling to sound model which also contains both direct connections from orthography to phonology, and a set mediated by hidden units. They characterize their model as a dual process model, with direct connections encoding regular or rule-governed relationships, while the hidden unit pathway mediates performance on exception words. When the hidden unit pathway is deleted, the present model performed quite well reading regular words and was extremely impaired at reading exception words.

This model does not have this property[3]. The semantic path is able to read many exception words, even in the absence of hidden units mediating orthography to phonology (see Harm, 1998, for elaboration). The Zorzi et al. (1998) model has no semantic representations, and is hence much more impaired at reading exception words in the absence of hidden units mediating orthography and phonology. Further, we utilized distributed phonological representations, unlike the localist phonological units used by Zorzi et al. Direct connections are much better able to encode regular spelling to sound relationships when localist units are used. The Zorzi et al. (1998) model without hidden units is worse at reading exceptions than ours because it lacks a semantic path, and better at reading regulars because it utilizes localist representations. Hence, while we agree with Zorzi and colleagues that direct connections facilitate the model's performance on nonword reading, we do not agree that their introduction corresponds to a qualitatively distinct second "process" which cleanly dissociates the types of words which the model can read.

Training Procedure. The model operated using a continuous form of recurrent backprop through time (Pearlmutter, 1989), modified slightly to accumulate the input to units rather than the output. Specifically, the input to the ith unit x_i at time t is defined as:

$$x_i(t) \quad = \quad (1-\alpha)x_i(t-1) + \alpha \sum_j^N w_{i,j} o_j(t-1) \tag{1}$$

while the output is $o_i(t) = f(x_i(t))$. In this way, input to the units ramps up over time, at a rate proportional to the time constant α. Time in the network was discretized over 12 samples, and run for 4 units of whole time, giving an integration constant $\alpha = 4/12 = 0.333$. Thus at each time step, the input to each unit becomes 0.333 times closer to what its environment is dictating. Note that in the case of $\alpha = 1.0$ this reduces to normal discrete backprop through time (as in Williams & Peng, 1990). Error is backpropagated through the network according to the normal backprop rules, except that the backprop term $\frac{\delta E}{\delta x_i}$ is gradually ramped up according to the same formula as Equation 1 (the dynamics of error propagation and activity propagation must be the same if training is to be stable).

A learning rate of 0.05 was used throughout training. Items were presented to the model according to a probability of presentation which was proportional to the square root of their frequency in the Wall Street Journal Corpus (Marcus et al., 1993). The frequencies were compressed using the square root to make simulations more computationally tractable: if probabilities of presentation were directly proportional to frequency, the lowest frequency item would be presented approximately once for every 23 million items. See Plaut et al. (1996) for discussion of the effects of various frequency compression schemes.

In Harm and Seidenberg (1999), the reading model was trained in two stages. In the first stage, the phonological attractor was trained on the phonological forms of words. This was to simulate the knowledge of the sound structure of language that children possess prior to training in literacy. This knowledge could then be used by the network when learning to map orthographic forms onto this pre-structured phonological attractor. For this model we have expanded this idea to include semantics: not only does the pre-literate child know about the sound structure of the target language, but also can map many of those phonological forms onto meaning and vice versa. The child also knows a good deal about the semantic structure of the world around them (for example, living things tend to have eyes, vehicles tend to have wheels, etc.) Much in the way that the phonological attractor explored by Harm and Seidenberg (1999) was intended to simulate implicit knowledge about the co-occurrence of phonological features in the target language, the semantic attractor used here is intended to simulate implicit knowledge about the co-occurrence of semantic features in the world.

To simulate these kinds of prior knowledge and their use in reading, the model was trained in two stages. In the first, pre-literate stage, the model was trained on 6103 items, with no orthographic information present. This is to train both the weights from phonology to semantics and back, and the phonological and semantic attractors. This training stage was modeled using two concurrent tasks: a speaking/listening task, and a hearing/thinking task. In the speaking/listening task, on 80% of the trials the model mapped semantics to phonology through the hidden units and phonological cleanup units. On the remaining 20% of the trials the phonological cleanup units were trained to retain a decaying pattern of activity in phonology (see Harm & Seidenberg, 1999, for a similar scheme used in reading). In the second, hearing/thinking task, an analogous training regime was used in reverse. On 80% of the trials, the model had to map phonological representations through the hidden units and semantic cleanup units onto a semantic representation. On 20% of the trials, the semantic cleanup weights were trained to retain a decaying semantic pattern. The speaking/listening and hearing/thinking tasks were trained separately for computational efficiency: because no weights are shared between the two tasks, two separate computers could be used simultaneously.

The hearing/thinking task was trained for 700,000 word presentations, at which point its performance had asymptoted. The speaking/listening task was trained for 500,000 word presentations, at which point its performance had reached asymptote. This concluded the first, pre-literate phase of training.

In the second stage of training, the orthographic part of the model was trained. The training set was expanded to 7455 items, in part to simulate the expanse of vocabulary that comes with the acquisition of literacy and in part to provide better coverage of the orthographic neighborhoods in English. Orthographic representations were clamped, and targets provided for both semantics and phonology. The model was able to make use of prior knowledge

---

[3] We tested the model on a set of exceptions from the "surface list" (Patterson & Hodges, 1992) and nonwords from the the graphemic complexity test of Howard and Best (1996). Training the model and then deleting the connections from the hidden units mediating orthography and phonology yields 67% performance on the nonwords and 51% performance on the exceptions, not the strong dissociation reported by Zorzi et al.

of the phonological and semantic structure of the target language learned in the first stage of training. It was trained in this manner for 650,000 word presentations.

Testing Procedure. For testing, the model was run for 24 samples over 8 units of whole time, again giving an integration constant $\alpha = 0.333$.

To measure latency, the model's output was sampled when the phonological output had settled. Formally, settling was defined as the point in time in which none of the model's phonological outputs had changed by more than 0.3 for four consecutive samples, or 1.333 units of whole time. This choice of parameters gave a reasonable range of latencies in word reading without pinning values at the high or low end. The amount of time which had elapsed was then recorded as the latency for the item. As is common in empirical studies, latencies were only recorded for items which the model produced the correct output.

The sum squared error was also computed from the model's outputs at the point of settling by taking the square of the difference between the model's phonological output for each feature and the target output. Sum squared errors were taken over all items whether the correct output was produced or not.

Phonological outputs were evaluated by finding the nearest phoneme, in euclidean distance, for each slot. The output was scored correct if all phonemes for all slots were correct. The semantic output of the model was considered correct if each the output feature was within 0.5 of its target value (in other words, unit outputs were rounded up or down to 0 or 1.0 and then tested for an exact match to the target).

To simulate the effects of acquired phonological impairments, the normal model was tested under conditions of noise within the phonological attractor. The noise took the form of multiplicative gaussian noise on the weights within the phonological attractor (the phonological units, and the phonological cleanup units). Noise with a standard deviation of 2.0 was used for the impaired condition. All damage was confined to the phonological attractor. This form of impairment is identical to the severe developmental phonological impairments used in Harm and Seidenberg (1999) to simulate the most severely impaired children from the Manis et al. (1996) study. Harm and Seidenberg (1999) explored several forms of phonological impairment, including the imposition of weight decay on the phonological connections, lesioning phonological cleanup units, and severing connections within the phonological apparatus in addition to using noise within the weights. Importantly, all forms of impairment to phonology led to a decrement in nonword reading over word reading, to varying degrees.

The imposition of noise within the phonological weights was chosen for these simulations because it had been shown to lead to the greatest degree of nonword impairment. This choice should not be taken as an literal model of the exact source of phonological impairment in either developmental or acquired phonological dyslexia, nor as a claim that the phonological impairments seen in developmental and acquired phonological dyslexia have the same underlying cause. Severely impaired phonological representations can result from a number of causes, both acquired and developmental. Rather, these simulations should be viewed instead as an exploration of the effect of such severely impaired representations on reading.

In the comparisons that follow, the normal model was damaged and tested twenty times, each damage study using a different random number seed for the gaussian noise. Multiple damage studies were conducted to ensure that any results obtained were not due to one particular, spurious damage pattern but rather would hold across a series of quantitatively different impairments.

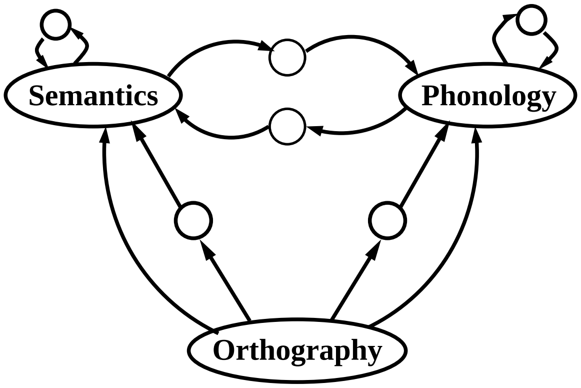

Figure 3. Implemented model used in Simulations 1 and 2.

Results

The items were tested on the normal model and twenty damage studies. The normal model and the impaired simulations yielded a pattern of results qualitatively similar to patient MJ and the normal undergraduates, as shown in Table 1.

Table 1
Effect of Graphemic Complexity

	Graphemic Complexity		
	Simple/3	Simple/5	Complex
MJ			
Accuracy	48%	36%	27%
Latency	626	611	730
Experiment 1			
Accuracy	99%	98%	93%
RT	549	598	633
Normal Model			
Accuracy	100%	95%	80%
Latency	5.90	6.26	7.42
SSE	0.00	0.13	0.48
Impaired Simulations			
Accuracy	64%	55%	35%
Latency	8.36	8.36	10.65
SSE	1.45	1.73	2.70

Note. All latencies are for correct items. SSE scores are for all items.

The sum squared error (SSE) of the normal model revealed an effect of complexity ($F(2,128) = 11.13$, $p < 0.001$). The impaired models also revealed an effect of complexity on SSE ($F(2,128) = 20.23$, $p < 0.001$). Combining the SSE scores for the normal and impaired models, an interaction between condition and impairment was reliable ($F(2,128) = 9.89$, $p < 0.001$). By inspection, the effect of the phonological impairment was to exaggerate the differences in sum squared error across conditions.

The latencies of the normal model were also subjected to an ANOVA. The effect of complexity on the sum squared error (SSE) was significant ($F(2,117) = 7.75$, $p < 0.001$). The effect of condition on the impaired models' latencies was also reliable ($F(2,127) = 19.65$, $p < 0.001$). The effect of condition on the impaired models' accuracies were reliable ($F(2,38) = 44.24$, $p < 0.001$).

Qualitatively, the latencies and SSE scores for the normal model match the RT scores of the subjects quite closely. The latencies for the impaired model match the general shape of MJ's latencies as well: the simple/3 and simple/5 conditions being quite close, and the complex items much higher.

Graphemic Complexity: Discussion

The point of the demonstrations thus far is simple. Orthographic manipulations affect how hard it is to generate phonological codes; nonwords that are harder for normal subjects are also harder for the normal model; phonological damage exaggerates this effect in both patient MJ and the phonologically impaired model. The results are consistent with the hypothesis that phonological dyslexia involves a phonological (but not orthographic) impairment.

2. Orthographic Modulation of the Pseudohomophone Effect: MJ

Howard and Best (1996) tested subject MJ for an effect of visual similarity on pseudohomophones. They created a set of pseudohomophones and matched nonwords (e.g., PHOCKS and SNOCKS). These were further split according to their visual similarity to the homophonous word; items such as GERL are similar to their homophone while PHOCKS is visually distant from its homophonous word.

They found that the manipulation of visual similarity affected MJ's performance on the pseudohomophones, such that similar items were much more accurately read than dissimilar pseudohomophones. No such effect of visual similarity was seen for the control nonwords. Howard and Best (1996) argue that because the visually similar and dissimilar pseudohomophones are phonologically equal (both sets are pseudohomophones), and the difference between them is purely orthographic, differing performance on these items indicates an impairment beyond phonology. Specifically they argue such results, like the graphemic complexity manipulation, indicate an impairment in graphemic analysis.

Howard and Best further argue that this result has implications for theories of word recognition, particularly accounts of why phonological dyslexics have been found to perform better on pseudohomophones than nonwords. Patterson et al. (1996) and Plaut et al. (1996) suggested that this effect can be explained by interactions between meaning and phonology; partial activation of phonology activates semantic representations of the word, which in turn bolster the phonological activations further. Howard and Best (1996) claim their finding of an orthographic effect on the pseudohomophone advantage casts doubt on this account. They assume that nonwords (whether pseudohomophones or not) will not activate semantics directly. Thus, the only information that can activate semantics is phonology, and hence orthographic effects are not predicted by this account. Howard and Best (1996) state (p. 916):

> This claim ... means that neither GERL (a visually similar pseudohomophone) nor PHOCKS (a visually dissimilar one) will produce any semantic activation. The only activation of semantics will be from the interaction between the phonological representation and semantics. In this interaction both pseudohomophones benefit equally (as both are pseudohomophones!). The prediction of the Plaut et al. account is clear: There should be a pseudohomophone advantage relative to control nonwords, but this effect will be independent of the visual similarity between the pseudohomophone and its real word homophone.

The logic of this argument follows, however, only if effects on the pseudohomophone advantage are truly from orthographic similarity to the base word rather than other confounded differences between stimuli[4]. Considering the examples given above, GERL is clearly a simpler nonword than PHOCKS. An obvious question then arises, which follows from the discussion of graphemic complexity above. Are the orthographic effects on the pseudohomophone advantage really due to orthographic similarity, or due to other aspects of the stimuli? If, as suggested by the GERL/PHOCKS example, the visually dissimilar pseudohomophones are actually more complex than the similar ones, then we obtain the same predictions as in the graphemic complexity section. First, the pattern of errors seen in MJ should be seen in the RTs of normal subjects. Second, the same unimpaired model used in the graphemic complexity section should show the same effects in its performance, and third, the imposition of phonological damage should exaggerate the effect.

We used the Howard and Best (1996) stimuli to see if the pattern of errors seen in MJ were reflected in the reaction times of normal subjects. We then tested the stimuli on the same normal and impaired models used in

---

[4] The argument also hinges on the assumption that nonwords never activate semantics directly, but only via phonology. This assumption will be addressed in the General Discussion.

the graphemic complexity investigation to see if again the model could link the behavior of normals and impaired populations through the addition of a phonological impairment. Finally, we will present an analysis of the stimuli which demonstrates the presence of a confounding factor that contributed to MJ's pattern of errors.

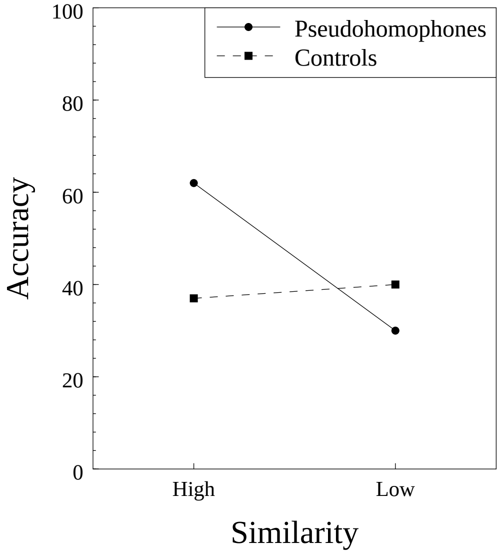

Figure 4. MJ accuracy: visual similarity by pseudohomophony interaction (from Howard and Best 1996).

Experiment 2: MJ Visual Similarity by Pseudohomophony

The design of this study was similar to Seidenberg, Petersen, MacDonald, and Plaut (1996), in which subjects were asked to name nonwords in three conditions: immediately, and after medium and long delays. They found that a pseudohomophone advantage was present even at the long delay. They argued that this suggests an articulatory basis for the pseudohomophone advantage: familiar phonological patterns are articulated more easily. The same paradigm was used in this study, in which items were named by subjects under immediate or delayed conditions. The prediction was that in the immediate condition, the subjects' pattern of RTs would match that of subject MJ. In delayed conditions, only a pseudohomophone advantage would be seen.

Method

Subjects. The subjects were USC undergraduates, as in Experiment 1.

Stimuli. The items for this experiment were taken from Howard and Best's (1996, p. 920) experiment. A total of 51 high similarity pseudohomophones and their matched controls, and 49 low similarity pseudohomophones and their controls were included in the list; a total of 40 of the 240 items used by Howard and Best were removed because of dialect differences (e.g. words like GLARCE, which in American English is not homophonous with GLASS). Items were counterbalanced by condition, with the added constraint that an item and its matched control never appeared within 5 items of each other. Six items and their matched controls were removed from consideration in reaction time analysis, because less than 50% of the subjects provided the correct response (e.g. failing to read YIGHND to rhyme with FIND). This left 47 items in each condition.

Table 2
Visual Similarity by Pseudohomophony Summary

	Pseudo-homophone		Control	
Condition	High	Low	High	Low
MJ				
Accuracy	62%	30%	37%	40%
RT	738	886	819	725
Normal Subjects				
Short RT	614	646	632	633
Medium RT	393	404	415	420
Long RT	365	376	383	384
Stimuli Analysis				
Mean $\sum \log f$	549	223	456	458
Mean $\sum$ Types	221	91	191	185
Normal Model				
Accuracy	98%	81%	89%	89%
SSE	0.24	0.51	0.30	0.29
Impaired Models (average)				
Accuracy	56%	40%	46%	50%
SSE	1.76	2.47	2.08	2.00
Damage Study #7				
Accuracy	59%	34%	41%	38%
SSE	2.11	2.66	2.59	2.48

Note. See text for a definition of $\sum \log f$ and $\sum$ Types.

Procedure. For each trial, a fixation point appeared on the screen for 750ms. This disappeared for 100ms, then the item to be named was presented. Subjects were instructed to pronounce the item to be named only when brackets appeared around the item. For each subject, a random one third of the items were presented with the brackets; the subject was instructed to name these items immediately. Another random one third of the items were presented with a medium delay; the brackets appeared 1,000 ms after the item appeared. The remaining items were presented with the bracket appearing 1,500 ms after the item appeared; this is the long delay condition. Immediate, medium and long delay conditions were all intermixed so the subjects could not know which condition would be seen on any given trial.

Results

All responses shorter than 100ms or longer than 2,000ms were coded as equipment failures and deleted from consideration. All remaining responses were trimmed by delay condition: each RT greater or less than 2.5 standard

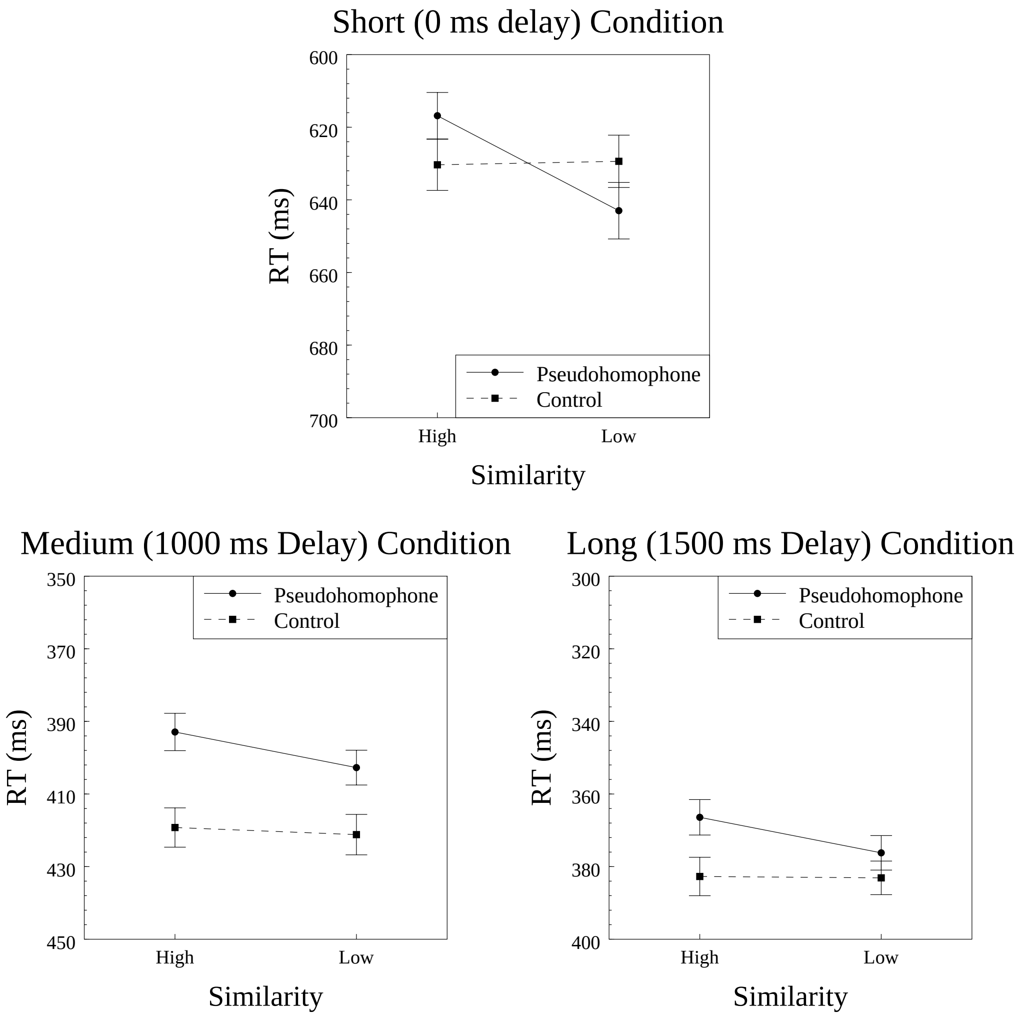

Figure 5. Experiment 2: latencies for immediate naming, and medium and long delay.

deviations from the mean for the delay cell was deleted.

The latency data from all three conditions are plotted in Figure 5, and summarized in Table 2. For the immediate naming condition, analysis of the reaction time data revealed a reliable interaction of pseudohomophony and visual similarity by subjects ($F1(1,24) = 5.3$, $p < 0.03$) and a marginally significant interaction by items ($F2(1,184) = 3.0$, $p < 0.087$), qualitatively replicating the naming performance of MJ (compare rows in Table 2, and to Figure 4, replotted from Howard and Best (1996), Figure 8c(ii)).

For the medium condition (Table 2), there was a reliable effect of pseudohomophony by subjects ($F1(1,24) = 3.89$, $p < 0.002$) and by items ($F2(1,184) = 9.47$, $p < 0.002$). There was a marginal effect of visual similarity by subjects ($F2(1,24) = 3.89$, $p < 0.06$) but not by items ($F2 < 1.0$). An interaction of pseudohomophony and visual similarity did not approach significance by subjects or items ($F < 1$).

The long condition, again shown in Table 2 was qualitatively similar to the medium condition. Again there was a reliable effect of pseudohomophony by subjects ($F1(1,24) = 5.9$, $p < 0.02$) and by items ($F2(1,184) = 4.8$, $p < 0.03$). An effect of visual similarity approached significance by subjects ($F1(1,24) = 3.3$,p<0.08) but not by items ($F2(1,184) = 1.2$, $p < 0.28$). Again there was no interaction between the two.

Combining the data from the three conditions, the interaction of pseudohomophony and visual similarity is marginally significant by subjects ($F(1,24) = 4.1$, $p < 0.053$) but does not approach significance by items ($F(1,184) =< 1$). The main effect of pseudohomophony was significant by subjects ($F(1,24) = 20.1$, $p < 0.001$) and by items ($F(1,24) = 5.6$, $p < 0.02$), while the main effect of visual similarity was significant by subjects ($F(1,24) = 16.6$, $p < 0.001$) but not by items ($F(1,184) = 1.7$, $p < 0.2$).

Experiment 2 demonstrates that normal subjects show a pattern of reaction times in immediate naming very similar to MJ's untimed accuracy (Figure 4). The medium and long

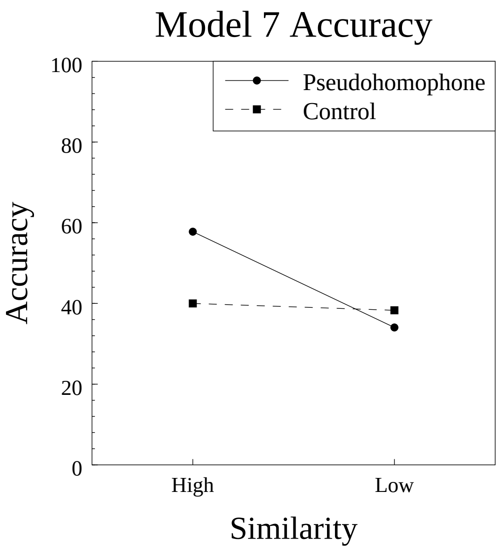

Figure 6. Simulation 2: mean accuracies for damage study #7. Compare with MJ's performance, depicted in Figure 4.

delay conditions demonstrate an advantage for pseudohomophones (mean 16ms) reflecting an articulatory benefit to pronouncing a familiar word form. This 16ms advantage is very close to the 13ms advantage reported by Seidenberg et al. (1996).

Simulation 2: MJ Visual Similarity by Pseudohomophony

To determine if a purely phonological impairment could yield a pattern of reading accuracies similar to that observed with subject MJ, a second simulation was undertaken. The plan was the same as for the first simulation: test the model in both the normal and impaired conditions to determine if the imposition of a phonological impairment exacerbated difficulties in the stimuli.

Method

Materials. The items used by Howard and Best (1996) (minus those excluded from the empirical study due to dialect differences) were tested using the same normal and impaired models from Simulation 1. Three pseudohomophones (ARTCH, HELLTH and LARDGE) and their matched controls were deleted from analysis because they contain graphemes in positions which the network had never been exposed to during training (see Harm & Seidenberg, 1999, for discussion). This left 182 items.

Procedure. The same normal and impaired simulations reported in Simulation 1 were used. The method of presenting items was identical as well.

Results

Table 2 summarizes the modeling results. The impaired model exhibits a reliable interaction of pseudohomophony and visual similarity on its accuracy scores ($F(1,19) = 67.19$, $p < 0.001$). Qualitatively, the mean accuracy scores for the impaired models match that of MJ; an effect of visual similarity on the pseudohomophones but not on control items. The magnitude of the effect of visual similarity on the pseudohomophones is somewhat lower for the models than for MJ. However, inspecting the performance of the twenty damage studies, one (the seventh) exhibited a pattern of errors very close in magnitude to MJ; this is shown in the final row of Table 2. The performance of this damage study is depicted in Figure 6. Performance for this "patient" on the graphemic complexity items was also very similar to MJ: simple/3 62%, simple/5 41%, and complex 22%.

The effect of similarity was reliable for the impaired model's sum squared error ($F(1,178) = 6.50$, $p < 0.02$). There was no reliable effect of sum squared error for the normal model, although the pattern qualitatively matches those of the subjects in the immediate condition, and of patient MJ's errors. There were no reliable effects on the normal or impaired models' latencies.

Combining the SSE scores for the normal model and the damage studies, a reliable interaction was obtained ($F(1,178) = 4.77$, $p < 0.03$). The effect of the phonological impairment is to cause a much larger spread in the effect of visual similarity on the pseudohomophones. Again, as was shown in Simulation 1, the phonological impairment exaggerated differences in the difficulty of the stimuli. A qualitative pattern in SSE in the normal model becomes a pattern of error in the phonologically impaired model.

Stimulus Analysis 1: MJ Visual Similarity by Pseudohomophony

What is the source of the effect of visual similarity in MJ's reading? One account of the pseudohomophone effect relies on interplay between the semantic and phonological regions of the triangle model (Patterson et al., 1996). The account of Seidenberg et al. (1996) also appeals to the triangle model, but places the locus of the effect in articulatory output. In either event, however, no effect of visual similarity is predicted; in both cases the predicted source of the effect takes place after the translation from spelling to sound. In contrast, Howard and Best argue, the DRC model does predict such an effect. Hence, the discovery of visual similarity modulation of the pseudohomophone effect could help mediate between the two models.

However, the predictions of the source of the pseudohomophone effect do not rule out the possibility that one can obtain an effect of visual similarity if one makes the visually dissimilar items more difficult to read along other dimensions. One can imagine that when constructing pseudohomophone pairs, one of which is orthographically similar to the root word and one that is distant, it is difficult to avoid introducing more complex, low frequency or unexpected letter combinations into the distant items. Such a tendency appears to be true of the Howard and Best items.

The low frequency letter string PH appears in 16 of their pseudohomophones; 15 of the low similarity items and one of the high similarity items. The onsets KW or SKW don't appear at the beginning of any English words, yet appear twice in the Howard and Best stimuli: both times as low similarity pseudohomophones.

To test for a bias toward such irregular letter strings, all 240 of the Howard and Best visual similarity by pseudohomophony test items were analyzed in terms of their orthographic onsets and rimes.

Method

Procedure. Onsets and rimes were measured rather than bigram frequencies because it was felt that they better capture the orthographic strangeness of a word. For example, the very odd word YIGHND has a summed positional bigram frequency count of 4606 as measured from the CELEX online dictionary of English monosyllables; this matches the bigram counts for very ordinary words like TRAMPS (4677), SKETCH (4705), and SWEEPS (4594). This is because YIGHND contains high frequency digrams such as ND in the 5th position (count: 3231), and this washes out the fact that HN virtually never occurs in the 4th position, and in general obscures the fact that IGHND is a truly difficult orthographic rime in English. Further, analyses of onsets and rimes are reasonable because the onset-rimes distinction is psychologically relevant in English, particularly for sounding out words (Treiman, 1986, 1992). An analysis of the frequencies of onsets and rimes in the Howard and Best (1996) study was undertaken.

The CELEX online dictionary of words and frequency counts was searched for monosyllabic words containing the onsets and rimes found in the Howard and Best stimuli. The total number of word types containing the onset was summed for each word. The calculation was repeated for the rimes. For example, the word PRIFE has the onset PR, which occurs in 106 word types in the CELEX monosyllables, so it gets an onset count of 106. The rime IFE occurs in 6 word types in the CELEX monosyllables, so it gets a rime score of 6. The logarithm of the frequencies of these words ($\sum \log f$) were also summed and recorded[5]. Taken together, these two measures are a measure of the typicality of a word's onset or rime.

Results

The summed onset types and $\sum \log f$ measure are plotted in Figure 7, and listed numerically in Table 2, along with MJ's accuracy and latency data. For the number of onset types, there is a reliable interaction between pseudohomophony and visual similarity ($F(1,236) = 19.2$, $p < 0.001$). The nature of the interaction is that visual similarity has a significant effect on the pseudohomophones ($F(1,118) = 41$, $p < 0.001$) but not the controls ($F < 1$). For the onset $\sum \log f$ measure, the effect is qualitatively the same: an interaction of visual similarity and pseudohomophony ($F(1,236) = 20.7$, $p < 0.001$) such that visual similarity is significant for the pseudohomophones ($F(1,118) = 38.4$, $p < 0.001$) but not for the control items ($F < 1$). The only effect seen in the rimes was a main effect of visual similarity ($F(1,236) = 6.1$, $p < 0.015$) on the number of types. Compare with the data for MJ depicted in Figure 4.

The onset analyses of the MJ stimuli do not match the subject data well in one regard. There is a tendency for the subject to perform better on the pseudohomophones than predicted by the onset analyses. Recall, however, the reaction time advantage in pseudohomophone naming for normal subjects found in this study and Seidenberg et al. (1996). This advantage was present even in delayed naming, suggesting a motoric/articulatory locus for the effect. It is very possible that the mismatch between the onset analyses and the subject data is due to this factor, which the stimuli analysis would not be sensitive to.

Visual Similarity: Discussion

Howard and Best (1996) described a subject whose nonword reading accuracy was dependent on pseudohomophony and visual similarity. They argue that this contradicts the predictions of the Plaut et al. model. Further, the effect of visual factors on the patient's naming has been argued to indicate a non-phonological basis for the patient's impairment (Coltheart, 1996; Howard & Best, 1996).

We have shown that MJ's accuracy scores can be accounted for by a simulation having only a phonological impairment. The reaction times of normal subjects corresponds to the behavior of the unimpaired simulation, and the imposition of phonological damage to the model creates a pattern of errors very similar to MJ.

One potential source of these effects is a confounding factor unintentionally present in the Howard and Best stimuli: the type and token frequency of words containing the onset of the presented nonwords. Simply put, patient MJ, the subjects and the model all find words beginning with KL more difficult than words beginning with BL, because few words begin with KL. Analysis of the model's error under conditions of purely phonological impairment reveals that such impairment not only preserves but exacerbates this difference, suggesting a clear line between the normal subject's RT data and patient MJ's accuracy data. The accuracy results of Simulation 7 confirm that a solely phonological impairment could dramatically affect nonwords with rare onsets much more so than those with more common ones. It is therefore quite plausible that this stimuli bias we have isolated (onset typicality) drove all the visual similarity effects cited by Howard and Best (1996).

---

[5] Log compression was used because it has been found in word recognition research that the log of the frequency of a word is a much better correlate of naming time than raw frequency. Such compression has the effect of preventing very high frequency items from masking variation in lower frequency items.

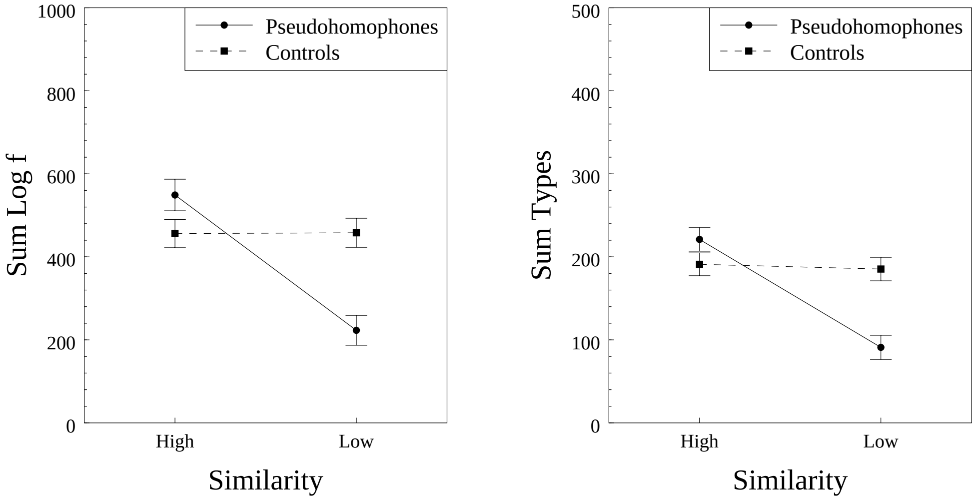

Figure 7. Stimuli analysis of visual similarity by pseudohomophony effect reported by Howard and Best (1996).

3. The Case of Patient LB

Patient LB (Derouesné & Beauvois, 1985) is argued to provide further evidence against the phonological impairment hypothesis (Coltheart, 1996). First, LB exhibits a visual similarity by pseudohomophony interaction similar to that reported by Howard and Best (1996) for subject MJ. Second, although LB exhibited impaired nonword reading, there was no strong evidence of a phonological impairment. Coltheart (1996) argues that LB thus provides at least once case in which the signature deficits of phonological dyslexia can be found with no phonological deficit. These claims will be considered in turn.

Derouesné and Beauvois (1985) do not report matching their items for orthographic properties, and inspection of their items suggest problems in the interpretation of LB's responses. Many of the visually similar pseudohomophones are highly similar to the target word, some differing only by the omission of one repeated letter (e.g., OCUPÉ for OCCUPÉ, ARIVÉ for ARRIVÉ, and COFRÉ for COFFRÉ). Conversely, many of the visually dissimilar pseudohomophones are orthographically unusual or illegal (e.g., SYVYL for CIVIL). The Derouesné and Beauvois (1985) items were therefore analyzed more closely.

Derouesné and Beauvois (1985) constructed a 2x2x2 design in which pseudohomophony, visual similarity and complexity were crossed. For the purposes of this analysis the data are collapsed across the complexity manipulation. Table 3 shows LB's accuracy. As with MJ, LB demonstrated an effect of visual similarity on the pseudohomophones but not on the control nonwords.

To obtain an estimate of French onset and rime frequencies, a set of 2.3 million words were obtained from electronic transcripts of the Canadian Parliament. As with the English analysis of MJ's stimuli, the onsets and rimes were coded for log frequency and number of word type containing that onset. Items with no onset were coded as zero. The same method used in Stimuli Analysis 1 for calculating onset and rime characteristics (number of types and $\sum \log f$) was used for this analysis.

No reliable effects were found for the rimes. For the $\sum \log f$ measure of the onsets, there was a reliable effect of similarity ($F(1,236) = 3.9$, $p < 0.050$) but no interaction ($F < 1$). However, inspection of the items shows a strong tendency in the direction of the effect; the lack of a reliable effect on the means is difficult to interpret because there are large outliers. For example, in the low similarity pseudohomophone condition, the mean $\sum \log f$ is 1029, but the standard deviation is quite large (1534). The values range from 0 to 4564, with 65% of the items numbering 102 or less, and no items in the range of 102 to 1009. The distribution is thus highly bimodal, with the mean being determined largely by the high and low extremes. The high similarity pseudohomophone condition, in contrast, has 41% of its items numbering 102 or less, and 26% of the items distributed between 102 and 1009. The large qualitative difference between the distribution of these items cannot be seen by comparing the means, where one extra very high item masks a very skewed distribution of items at the low end. When the median values are considered, the results map much more closely onto LB's performance. Table 3 shows the median values by cell. For the number of types analysis, there was a marginal effect of similarity ($F(1,236) = 3.8$, $p < 0.051$) and no interaction ($F < 1$).

The median onset counts shown in Table 3 match LB's accuracy quite well. There is a very small effect of visual similarity on LB's the nonword reading; this effect is seen in the onset analysis. A much greater effect of similarity is seen on the pseudohomophones.

Patient LB (Derouesné & Beauvois, 1985) was also said to have normal or near-normal word reading abilities, but a

Table 3
LB Stimuli Analysis

	Pseudo-homophone		Control	
Condition	High	Low	High	Low
LB				
Accuracy	85%	52%	35%	25%
Stimuli (Median)				
$\sum \log f$	1204	556	612	551
$\sum$ Types	763	477	430	348
Stimuli (Mean)				
$\sum \log f$	1205	1029	1104	566
$\sum$ Types	795	689	739	371

severe impairment reading nonwords. Interestingly, unlike every other reported case of phonological dyslexia over the past 15 years, no phonological impairments were reported for LB (Coltheart, 1996). LB's case, then, was taken as evidence by Coltheart that a phonological impairment need not underlie all cases of phonological dyslexia.

Derouesné and Beauvois (1985) tested LB's phonological skill on a several tasks.

1. Nonword phonemic segmentation. LB was required to produce the last phoneme in a spoken word. He scored 40/40 correct.
2. Nonword blending. Three phonemes were pronounced to LB; he had to combine them into a word. He scored 27/30 correct.
3. Word-in-word detection. Two pictures were shown to LB; he had to decide if the sound form of one word represented by one picture was contained within the sound form of the word represented by the other picture. He scored 36/40 correct.

Given these results, Coltheart (1996) concluded (p. 755) "hence there is no evidence of a phonological impairment". He further stated that LB read words with "reasonable accuracy", ranging from 74% to 98% for various types of words. However, LB's score on reading simple nonwords (two letter monosyllables) was 34/40, or 85%, while his reading of more complex nonwords was 19/40 (48%). His ability to assemble simple three phoneme sounds was 27/30, or 90%.

One could question the labeling of a nonword impairment as "severe" based on the more difficult items (48% correct), and simultaneously labeling his phonological skills as "normal" based on the simplest items (on which he scored 90%). Derouesné and Beauvois (1985) do not provide performance statistics for normal subjects on their items in the phonological tests, but we can probably assume that normals would score very close to 100%. If these scores do indeed show "no evidence of a phonological impairment", then it is left unexplained why his scores were in fact not at ceiling. As observed by Patterson (in press), LB's performance on nonwords was actually high relative to many other phonological dyslexics that have been studied in the literature, and it is therefore not surprising that his phonological performance would be higher than many phonological dyslexics.

A further complication in interpreting LB's performance is the fact that he underwent a four year remediation program after the CVA, which attempted to restore his reading ability. As noted by Derouesné and Beauvois, the rehabilitation was based on attempting to train phonological reading. It is unclear exactly what the four year remediation program consisted of, but in many cases such phonologically based remediation includes intensive instruction in "phonics" skills. These skills involve things like phoneme blending, segmenting, matching and deletion; exactly the kinds of skills that Derouesné and Beauvois used years later to test his phonological skills. Improvements in such skills do not necessarily translate into improvements in nonword reading (Adams, 1990). The remediation program essentially failed to improve LB's reading, but may have made him very good at finding alternate strategies to blend and segment words. In short, the remediation program may have made him into an expert at generic phonological tests while still having very poor phonological representations.

A way to test this hypothesis would be to give LB a more sensitive test of phonological skill, such as in Werker and Tees (1987), where reading impaired children were tested on a categorical perception task. If LB showed a more flat categorical perception curve, like the reading impaired children studied by Werker and Tees (1987), then it could not be argued that he did not possess a phonological impairment, and we would be forced to conclude that the phonological tasks LB was assessed on originally were too confounded with the remediation program he had been subjected to too provide a true measure of his actual phonological skill.

It is argued by Coltheart (1996) that LB reads words with "reasonable accuracy", scoring from 74% correct to 98% correct. Derouesné and Beauvois, however, state (Derouesné & Beauvois, 1985, p. 406) that "The Alouette Reading Test … confirmed his serious reading impairment." LB's reading level on this standardized test placed him in the second year of elementary school. His reading was reported to be 183 words in 180 seconds, making 35 errors. Derouesné and Beauvois then tested LB on a set of words they constructed, and they found that his reading of these words, which are not reported to have not been normed or standardized in any way, ranged from 74% to 98%. They then concluded, as Coltheart does, that that "he could read words fairly well." (p. 410).

It is questionable whether reading at a second grade level, at about one word a second, really constitutes "fairly well", and "reasonable accuracy". Patterson (in press) argues that LB's word reading is in fact poor, and that he therefore does not constitute a "strong" dissociation, that is, his performance on one task is not qualitatively different from the other task (cf. Shallice, 1988, p. 227). Both word and nonword reading are fundamentally impaired. Patterson suggests that if we are to accept that a single case of a given dissociation is as theoretically important as a legion of contrary cases, we should at least require that the single case be a "strong" dissociation.

To summarize, it has been shown that the stimuli used to demonstrate LB's visual similarity by pseudohomophony interaction is at the very least suspect, and does not provide strong evidence of such an interaction. Further, we claim it has not been firmly established that LB was free of a phonological impairment. Even if he was free from a phonological impairment, it is clear that his reading impairment was not a selective impairment in nonword reading. Therefore, the case of patient LB does not provide convincing evidence against the phonological impairment hypothesis.

4. General Discussion

Orthographic Effects

We have demonstrated that normal undergraduates and an unimpaired simulation show patterns effects of graphemic complexity in latencies similar to MJ's pattern of errors. We have demonstrated that damage to this simulation can yield patterns in errors that closely match MJ's. Therefore, the claim that this pattern of errors necessarily implicates from an impairment in graphemic parsing cannot be true. It is important to note that we have not demonstrated that a graphemic processing impairment cannot produce the reported effects. Rather we have simply shown that such effects can arise from a purely phonological impairment.

The same considerations apply to the effects of visual similarity to the base word on phonological processing. The effects reported by Howard and Best (1996) reflect the confound between similarity to the base word and orthographic complexity. This demonstration does not mean that the kind of effect Howard and Best were seeking would not be obtained with better materials. In fact we think that such effects probably do occur for reasons that follow from the behavior of the triangle model. Looking at the model in Figure 2, there is a pathway from orthography to semantics which leads into phonology. It is often assumed that nonwords do not produce significant activation along this pathway, but this may be untrue. For example, nonwords that are highly similar to a single word (e.g., GARDIN, NERSE) may activate the word's semantics sufficiently to support pronunciation.

The simulations presented here provide support for this hypothesis. We tested the Howard and Best (1996) visual similarity by pseudohomophony items on the model after totally lesioning the pathway from orthography to phonology, so that pronunciations could only be generated semantically. We found that eight of the 120 pseudohomophones could be correctly pronounced: BLUD, DETT, FRUNT, GERL, KAMP, MOWTH, MUNTH, SLOMP. For example, FRUNT produced correct phonological output because no other items in the training set fit the orthographic pattern FR_NT. Thus FRUNT activated enough of the semantic features for FRONT to produce the correct phonological output. The items that the network could pronounce all have relatively few orthographic neighbors; we would not expect a pseudohomophone like KAR, which has multiple close neighbors, to produce reliable semantic activation. The item KAR is no closer to CAR than to BAR,FAR,JAR,PAR or TAR.

Importantly, the semantic activation for the pseudohomophones is generally lower than for words, so it is not the case that the network cannot distinguish FRUNT from FRONT; see Plaut (1997) for relevant simulations.

Other subword regularities can be inferred by the semantic route. Words ending in ED tend to activate the <past tense> feature; this semantic feature in turn tends to activate a /d/ or /t/ in the final phoneme. For the 16 nonwords and pseudohomophones ending in ED (such as PESSED, KOSSED), the network produces a /d/ or /t/ in the final phoneme on 12 of them.

We therefore do not want to argue that one could not find a demonstration of the interaction of visual similarity and pseudohomophony. The main point is that the existing studies do not show such an effect. Given the simulation results discussed above, we would expect not a visual similarity by pseudohomophony interaction, but rather a neighborhood density by pseudohomophony interaction. The relevant variable is not just closeness to the target word, but closeness to the target word and distance from other words. In any event, such a demonstration would not adjugate between models of word recognition, nor would it provide evidence of a graphemic processing impairment.

We conclude that the reported effects of graphemic complexity and orthographic modulation of the pseudohomophone effect do not provide evidence against the phonological impairment hypothesis.

Must it be Phonology?

It has been argued that a phonological impairment will necessarily impair nonword reading more so than word reading. However, is a phonological impairment the only way to get a nonword impairment? Both the triangle model (Figure 2) and the DRC model (Coltheart et al., 1993) have two pathways to phonology from print. The SM89 model allows for "semantic" reading, and the DRC model allows for "lexical" reading. In principle, therefore, both models allow for a form of acquired phonological dyslexia through disruptions in the direct pathway from print to sound. This requires a developed and intact lexical/semantic pathway. In the developmental simulations, however, disruptions of the orth→phon pathway affected exception word reading more so than nonword reading (Harm & Seidenberg, 1999).

Conclusions

Our results call into question the claim that phonological dyslexia is associated with orthographic deficits. Words vary in terms of orthographic complexity, which affects the difficulty of the orthography → phonology mapping. Phonological impairment magnifies these differences between stimuli. Hence effects such as those seen in MJ are also observed in normal college students and in models that have no orthographic impairment. Thus the patient data are consistent with the phonological deficit hypothesis and do not demonstrate an orthographic impairment. The analysis of the LB data suggest that it is important to carefully assess the phonological and reading abilities of patients before concluding that particular capacities are impaired or preserved. As we have demonstrated elsewhere (Harm & Seidenberg, 1999), a phonological impairment can be severe enough to affect performance on one task (e.g., nonword naming) but leave performance on other tasks (e.g., nonword repetition) unaffected. These issues need to be considered carefully in future research addressing the possible bases of reading and other (e.g., morphological; Joanisse & Seidenberg, 1999) impairments.

References

Adams, M. (1990). *Beginning to read.* Cambridge, MA: MIT Press.

Beauvois, M. F., & Derouesné, J. (1979). Phonological alexia: Three dissociations. *Journal of Neurology, Neurosurgery and Psychiatry*, *42*, 1115-1124.

Berndt, R. S., Haendiges, A. N., Mitchum, C. C., & Wayland, S. C. (1996). An investigation of nonlexical reading impairments. *Cognitive Neuropsychology*, *13*(6), 763-801.

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

Bishop, D. V. M. (1992). The underlying nature of specific language impairment. *Journal of Child Psychology and Psychiatry*, *33*(1), 3-66.

Castles, A., & Coltheart, M. (1993). Varieties of developmental dyslexia. *Cognition*, *47*(2), 149-180.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English.* New York: Harper & Row.

Coltheart, M. (1996). Phonological dyslexia: Past and future issues. *Cognitive Neuropsychology*, *13*(6), 749-762.

Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, *100*(4), 589-608.

Coltheart, M., Davelaar, E., Jonasson, K., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention & performance VI.* Hillsdale, NJ: Erlbaum.

Derouesné, J., & Beauvois, M. F. (1979). Phonological processing in reading: data from alexia. *Journal of Neurology, Neurosurgery and Psychiatry*, *42*, 1125-1132.

Derouesné, J., & Beauvois, M. F. (1985). The 'phonemic' state in the non-lexical reading process: Evidence from a case of phonological alexia. In K. Patterson, M. Coltheart, & J. C. Marshall (Eds.), *Surface dyslexia* (p. 399-457). Hillsdale, NJ: Erlbaum.

Farah, M. J., Stowe, R. M., & Levinson, K. L. (1996). Phonological dyslexia: Loss of a reading-specific component of the cognitive architecture? *Cognitive Neuropsychology*, *13*(6), 849-868.

Harm, M. W. (1998). *Division of labor in a computational model of visual word recognition.* Unpublished doctoral dissertation, University of Southern California, Los Angeles, CA.

Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*(3), 491-528.

Howard, D., & Best, W. (1996). Developmental phonological dyslexia: Real word reading can be completely normal. *Cognitive Neuropsychology*, *13*(6), 887-934.

Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Science*, *96*(13), 7592-7597.

Manis, F., Seidenberg, M., Doi, L., McBride-Chang, C., & Peterson, A. (1996). On the basis of two subtypes of developmental dyslexia. *Cognition*, *58*, 157-195.

Marcus, M., Santorini, B., & Marcinkiewicz, M. A. (1993). Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, *19*, 313-330.

Patterson, K. (in press). Phonological alexia: The case of the singing detective. In E. Funnell (Ed.), *Case studies in the neuropsychology of reading.* Erlbaum.

Patterson, K., & Hodges, J. R. (1992). Deterioration of word meaning: Implications for reading. *Neuropsychologia*, *30*(12), 1025-1040.

Patterson, K., & Marcel, A. J. (1992). Phonological ALEXIA or PHONOLOGICAL alexia? In J. Alegria, D. Holender, J. Junça de Morais, & M. Radeau (Eds.), *Analytic approaches to human cognition* (p. 259-274). New York: Elsevier.

Patterson, K., Suzuki, T., & Wydell, T. N. (1996). Interpreting a case of Japanese phonological alexia: The key is in phonology. *Cognitive Neuropsychology*, *13*, 803-822.

Patterson, K. E., Marshall, J. C., & Coltheart, M. (Eds.). (1985). *Surface dyslexia: Neuropsychological and cognitive studies of phonological reading.* London: Erlbaum.

Pearlmutter, B. A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation*, *1*(2), 263-269.

Plaut, D. C. (1997). Structure and function in the lexical system: Insights from distributed models of word reading and lexical decision. *Language and Cognitive Processes*, *12*, 765-805.

Plaut, D. C., McClelland, J. L., Seidenberg, M., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*(1), 56-115.

Sasanuma, S., Ito, H., Patterson, K., & Ito, T. (1996). Phonological alexia in Japanese: A case study. *Cognitive Neuropsychology*, *13*(6), 823-848.

Seidenberg, M. S. (1995). Visual word recognition: An overview. In P. Eimas & J. L. Miller (Eds.), *Handbook of perception and cognition: Language.* New York: Academic Press.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523-568.

Seidenberg, M. S., & McClelland, J. L. (1990). More words but still no lexicon: Reply to Besner et al. (1990). *Psychological Review*, *97*(3), 447-452.

Seidenberg, M. S., Petersen, A., MacDonald, M. C., & Plaut, D. C. (1996). Pseudohomophone effects and models of word recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *22*(1), 48-62.

Shallice, T. (1988). *From neuropsychology to mental structure.* Cambridge: Cambridge University Press.

Stanovich, K., Siegel, L., & Gottardo, A. (1997). Converging evidence for phonological and surface subtypes of reading disability. *Journal of Educational Psychology*, *89*(1), 114-127.

Treiman, R. (1986). The division between onsets and rimes in English syllables. *Journal of Memory and Language*, *25*, 476-491.

Treiman, R. (1992). The role of intrasyllabic units in learning to read and spell. In P. Gough, L. Ehri, & R. Treiman (Eds.), *Reading acquisition.* Hillsdale, NJ: Erlbaum.

Werker, J., & Tees, R. (1987). Speech perception in severely disabled and average reading children. *Canadian Journal of Psychology*, *41*(1), 48-61.

Williams, R. J., & Peng, J. (1990). An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, *2*, 490-501.

Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(4), 1131-1161.