

Modeling the Successes and Failures of Interventions for Disabled Readers

Michael W. Harm
Carnegie Mellon University

Bruce D. McCandliss
Sackler Institute

Mark S. Seidenberg
University of Wisconsin, Madison

A connectionist model of reading development previously used to simulate detailed aspects of developmental dyslexia (Harm & Seidenberg, 1999) was used to explore why certain classes of interventions designed to overcome reading impairments are more effective than others. Previous research has shown that interventions targeting the development of spelling-sound correspondences are more effective at promoting generalization skills than ones designed solely to increase phonological awareness. The simulations broadly replicate the patterns of success and failure found in the developmental literature, and provide explicit computational insights into exactly why the interventions that include training on spelling-sound regularities are more effective than those targeting phonological development alone.

The number of studies of interventions targeting developmental reading impairments is growing at a rapid rate (see Bus & Ijzendoorn, 1999, for a recent review). Such studies typically involve testing one or more interventions with children; outcome measures indicate the relative strengths and weaknesses of the different interventions being considered. The benefits of such work are obvious: they give direct evidence for which of the tested interventions provide a more effective use of time and resources for the particular population of children studied. The difficulty is that such studies are very demanding to conduct, and generally involve researchers working with children in school environments, removing children from their normal schoolwork, and obtaining consent materials from parents and school officials. Further, the studies often examine in-

terventions that differ along many dimensions, making it difficult to isolate the causal factors that determine the specific patterns of outcomes. It is typically impractical to perform a more controlled study in which two methods vary only by one cognitive factor. Therefore, when one intervention is found to work better than another, it is often not obvious exactly why this would be the case, making it difficult to draw broader conclusions from the results of a particular study. Such considerations suggest the need to establish explicit links between cognitive models of reading development and the insights one can draw from particular intervention studies.

As intervention studies continue to increase in frequency and specificity, developments are also underway in the literature on cognitive models of reading and reading acquisition that help to draw specific links between empirical studies of reading intervention and theoretical constructs within reading models. Researchers often use computational models to make contributions to our understanding of reading processes by implementing specific hypotheses about the architecture and mechanisms of the reading system, and analyzing the performance of the model to account for various empirical phenomena. Previous connectionist models of reading have simulated detailed aspects of reading performance in adults and impaired populations using general computational principles concerning learning, knowledge representation and processing. Such studies have typically focused on understanding the end-state of skilled reading, by constructing explicit computational models of some facet of reading in adults, such as on-line performance data, (e.g., Seidenberg & McClelland,

Michael W. Harm, Center for the Neural Basis of Cognition, Carnegie Mellon University; Bruce McCandliss, Sackler Institute for Developmental Psychobiology, Weill Medical College of Cornell University; Mark S. Seidenberg, Department of Psychology, University of Wisconsin, Madison.

This work was supported by a National Research Service Award from the National Institute on Deafness and Other Communication Disorders, F32-DC00425, the J. S. McDonnell Foundation's Cognitive Studies in Education (grant #98-3), NICHD grant HD29891, NIMH grant MH04756 and by a Research Scientist Development Award, NIMH K2-01188, to Mark Seidenberg.

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.

1989), acquired reading disorders resulting from brain damage such as deep dyslexia (Hinton & Shallice, 1991; Plaut & Shallice, 1993) and surface dyslexia (Plaut, McClelland, Seidenberg, & Patterson, 1996). Each of these models of adult reading utilized pre-specified phonological representations, precluding the possibility of gaining insight into one of the central findings in reading acquisition research: the impact of poorly developed phonological representations on reading development. Coltheart, Rastle, Perry, Langdon, and Ziegler (2001) have also presented a computational model of many aspects of word reading, but it also does not address how phonological information is acquired or shaped through exposure to print.

One promising approach that may aid in establishing more explicit and generalizable links between reading intervention research and cognitive models is to simulate more detailed aspects of the learning process with connectionist models, particularly with respect to factors known to be relevant to reading development (e.g., phonological awareness, and the composition of learning materials). Harm and Seidenberg (1999) constructed a simulation of the development of reading skill, and impairments that can arise during development. This simulation departed from previous work by including a phonological system that learns the sound structure of the target language prior to learning to read. The phonological system used by Harm and Seidenberg (1999) could be impaired in its development, giving different instantiations of the simulations different levels of phonological skill. Harm and Seidenberg (1999) explored the effect of different levels of phonological skill on reading acquisition, and related the results to the findings from an empirical study of disabled readers by Manis, Seidenberg, Doi, McBride-Chang, and Peterson (1996). More recently, models have been trained using actual materials from different beginning reading basals, to examine their relative effectiveness in the face of reading disabilities (Foorman, Perfetti, Seidenberg, Francis, & Harm, 2001).

Computational simulations of reading have now reached a level of development that allows them to be applied to questions concerning the effectiveness of interventions for disabled readers. The strengths of such simulations complement empirical studies of reading interventions: it is easy to test multiple hypotheses rapidly and under a wide set of conditions. The internal details of such models can be analyzed to observe direct causal effects without real-world confounds that can make it difficult to interpret behavioral studies. By tightly linking the training conditions, intervention method and evaluation metrics of a computational model to established empirical studies, analysis of the model can provide leverage for understanding observed behavioral effects.

In this paper, we present a first attempt to use a connectionist simulation of reading development to explain results found in the empirical literature on reading interventions. We focus on understanding the precise nature in which the

quality of phonological representations influences the way the model learns to read words, and how different interventions might influence the development of reading skill for an individual that starts out with impairments to phonology.

The Phonological Impairment Hypothesis

There is now extensive evidence that phonological impairments are a major factor related to developmental reading difficulties in children (National Reading Panel, 2000). Childrens' performance on tasks involving the manipulation of component phonemes in a word is a reliable predictor of reading skill (e.g., Bradley & Bryant, 1983; Lundberg, Olofsson, & Wall, 1980; Mann, 1984; Lundberg et al., 1980; Rosner & Simon, 1971). This has led researchers to conclude that there is a potentially causal link between phonological awareness skill and early literacy acquisition (e.g., Tunmer & Nesdale, 1985; Wagner & Torgesen, 1987). At the core of the phonological impairment hypothesis is the notion that phonological awareness skills may act as a critical factor for the success of decoding attempts (Liberman, Shankweiler, Fisher, & Carter, 1974) and might thereby have a profound effect on literacy acquisition. Other studies have shown that further refinement in phonological skills results from exposure to an alphabetic script (Morais, Cary, Alegria, & Bertelson, 1979; Read, Yun-Fei, Hong-Yin, & Bao-Qing, 1987), suggesting a reciprocal relationship between phonological skills and literacy (Morais, Alegria, & Content, 1987; Wagner, Torgesen, & Rashotte, 1994; see Harm & Seidenberg, 1999 for computational simulations of this two-way relationship).

A straightforward inference from the phonological impairment hypothesis is that remediating phonological awareness impairments should produce a direct positive impact on reading skills. Numerous controlled intervention studies have addressed this prediction in low-achieving pre-readers (e.g., Bradley & Bryant, 1983; Ball & Blachman, 1991; Blachman, Ball, Black, & Tangel, 1994) and at-risk children with demonstrated reading difficulties (Lundberg, 1994; Schneider, Ennemoser, Roth, & Kaspert, 1999; Schneider, Roth, & Ennemoser, 2000; Byrne & Fielding, 1995; Hatcher, Hulme, & Ellis, 1994). These studies provide support for the general claim that phonological awareness training can have a significant impact on phonological awareness skills and also impact on reading skills. However, such conclusions are somewhat limited by the presence of negative findings (i.e. Weiner, 1994) and mixed evidence concerning long-term follow-up results (for discussion see Bradley & Bryant, 1983; Olson, in press). Given the large number of such intervention studies, reliable conclusions are perhaps best drawn from quantitative meta-analyses (Bus & Ijzendoorn, 1999; Ehri & Nunes, in press) that quantify average effect sizes and provide consideration of relationships between different forms of intervention

programs and the magnitude of outcome effects. Considering over 20 controlled studies in the U.S., Bus and Ijzendoorn (1999) demonstrated that, relative to control, phonological awareness training had a medium to strong ($d=.73$) impact on phonological awareness skills, and a medium to strong ($d=.70$) impact on reading skills, providing direct support for the claim that phonological awareness training has a causal impact on reading skills. Similar, yet weaker, results were obtained with a more inclusive, yet less homogeneous, sample of 34 studies (Bus & Ijzendoorn, 1999). A more recent meta-analysis (Ehri & Nunes, in press) expanded this sample to 52 studies, and also demonstrated that phonological awareness training had a large impact on phonological awareness skills ($d=.86$) and a medium impact on reading skills ($d=.53$).

These findings support the often cited causal link between phonological awareness skills and reading ability upon which the phonological impairment hypothesis is grounded.

Puzzling Results

Given this support for the phonological impairment hypothesis, it might seem that the most effective way to address reading difficulties is to remediate the underlying phonological impairment directly. For example, one could engage children in the same speech activities that are used to assess these skills, such as activities requiring children to segment and manipulate speech sounds within words (i.e. Lundberg, Frost, & Petersen, 1988). Meta analyses by Bus and Ijzendoorn (1999) however, demonstrate that such speech-only approaches are minimally effective at impacting reading abilities ($d=.18$), and consistently worse than other variants that spend less time on speech based activities and more time on activities involving reading ($d=.88$) or letters ($d=.66$). Similar meta analysis results were obtained by Ehri and Nunes (in press) in which the reading improvement effect size from phonological awareness training with letters ($d=.67$) was roughly twice that obtained from similar speech-only activities ($d=.38$). These results hardly follow from the simple form of the phonological impairment hypothesis, and seem to require additional assumptions. Several possibilities have been advanced, including the notion that the presence of letters might serve to perceptually anchor perceptually elusive phoneme sounds (Adams, Treiman, & Pressley, 1998), and the notion that specific training in grapheme-phoneme associations may directly impact reading abilities (i.e. Ehri & Nunes, in press).

A second result from the intervention literature that doesn't seem to follow from the phonological impairment hypothesis is that the degree to which phonological awareness training leads to improvements in reading abilities is critically dependent on the timing of this intervention in relation to the onset of reading. Meta analyses by Bus and Ijzendoorn (1999) report that preschoolers show significantly greater benefits than do kindergarten or primary

school children. Similar results were obtained by Ehri and Nunes (in press), demonstrating greater reading improvements for preschoolers ($d=1.25$) than children in kindergarten ($d=.48$) or primary school ($d=.49$).

Why is it that the effectiveness of phonological awareness training on reading drops off so precipitously just after the onset of attempts at reading? Simple notions of 'earlier is better' seem to fail to account for such a dramatic shift in effectiveness precisely as reading attempts begin. Instead, it seems the phonological impairment hypothesis might need to be elaborated upon to account for these phenomena. Furthermore, phonological awareness activities might need to be elaborated to attain the same effectiveness for children who have begun to read as for preschoolers who have not.

These two factors mediating the impact of phonological awareness training on reading skills—whether oral language activities are integrated with print activities and whether interventions occur early or late in relation to the onset of reading—might interact in important ways. Some preliminary support for this possibility comes from a study that compared normal kindergarten children who had received no prior reading instruction with first grade children who had. Cunningham (1990) used these two groups to examine the impact of enhancing oral phonological awareness activities with instruction on “metalevel knowledge of when, where, how and why to use phonemic awareness within the reading context (p. 431).” In this enhancement, “the utility of the skill for reading activities was demonstrated and practiced (p.436).” The relative magnitude of reading benefits across these groups of children produced an interesting interaction between age and training method. The presence or absence of the enhancement that explicitly integrated phonological awareness with reading skills had a significantly larger impact for first grade students than for kindergarten students. Perhaps once reading experience begins, oral phonological awareness training is no longer as effective at driving reading skill improvements, and such activities must be enhanced or replaced by methods that stress the connection between these activities and the structure within printed words.

Perhaps another class of puzzling results comes from intervention studies that do not explicitly target phonological skills, but instead directly target decoding skills. Interventions that train phonological awareness through methods that strongly emphasize orthographic to phonological mappings have been demonstrated to be effective at enhancing reading abilities of early readers selected for reading impairments (National Reading Panel, 2000). For example, a study by (McCandliss et al., in press) used the “Word Building” technique based on the work of Beck (1989), which emphasized the role of letter-sound correspondences in the context of words. This study examined the impact of 20 sessions of a print-based decoding skills intervention on the reading skills of children who demonstrated reading impairments persisting after the first grade. The overlap

in both sound and spelling between words was emphasized by teaching words from small dense orthographic neighborhoods, changing successive words just one letter at a time. McCandliss et al. (in press) theorized that this intervention's technique of building each new word by changing a single letter in the previous word places pressure on the orthographic \rightarrow phonological (hereafter orth \rightarrow phon) system to form more componential mappings; that is, ones more sensitive to the internal parts or components of words. As each new word was formed by changing a letter from the previous word, children made decoding attempts. When attempts failed, the child was instructed to segment the word into letter sounds and blend them together. Complexity of word forms gradually increased over approximately 70 lessons, and children only progressed to more difficult material after surpassing an accuracy criterion on each lesson.

The McCandliss et al. (in press) study found that the Word Building intervention led to substantial improvements in phonological awareness and still larger improvements in decoding abilities. Presumably, the intervention activities improved children's nonword reading abilities by helping them to develop orth \rightarrow phon representations that were more componential. These results are consistent with other studies that focus intervention activities explicitly on decoding rather than phonological awareness skills (National Reading Panel, 2000).

Given the extensive evidence that phonological impairments can cause poor reading, the limitations of efforts to train phonological awareness through speech-related interventions that do not include an orthographic component are puzzling. If poor phonological skill, as measured by speech tasks such as phoneme deletion and syllable segmentation causes poor reading, then why do interventions that specifically target those skills have much less of an impact than those including training of spelling-sound correspondences? And why does the relative importance of augmenting phonological awareness training with activities that target spelling-sound correspondences increase as the child becomes more experienced with reading?

We draw upon the insights that came from using computational models to analyze the effects of phonological impairments on reading acquisition (Harm & Seidenberg, 1999), and the design principles, materials, and results of the McCandliss et al. (in press) intervention to present here a theoretical account of this puzzling result, which we have termed the "mapping hypothesis." We then present a connectionist simulation of a purely phonological intervention, and a simulation of the McCandliss et al. (in press) intervention.¹

¹Of the large array of possible interventions to simulate, we chose the McCandliss et al. study based on several considerations that made it very amenable to this modeling effort. The study utilized a fixed set of monosyllabic word materials for all participants. A small number of explicitly specified algorithmic

The Mapping Hypothesis

The central claim of the mapping hypothesis is that poor nonword reading is caused by the formation of non-componential, holistic representations in the mapping from orthography to phonology. A componential representation shares structure with other items, and hence pronunciation can be aided through this overlapping structure; non-componential representations lead to poor generalization because the overlap in structure is less apparent. Such non-componential representations are themselves caused by poor phonological representations at the beginning of the acquisition of reading. Further, the character of subsequent learning is dependent on existing representations; once poor representations have been formed, subsequent learning exploits the (poor) characteristics of these existing representations. This leads to two predictions:

1. Phonological awareness remediations that ignore the orth \rightarrow phon mapping system result in limited improvements, particularly after poor representations have already been formed in the orth \rightarrow phon system, because they failed to place direct pressure on the system to change mappings that were already learned.
2. Mapping remediation (letter-sound mapping, Word Building) changes the orth \rightarrow phon mappings, making them more componential, and thus has a greater impact on older children, especially for generalization to novel words.

In this paper, we explore these predictions using a connectionist model of reading acquisition (described below; see also Harm & Seidenberg, 1999) which allows experimental manipulation of the phonological system and the orth \rightarrow phon system. We first recount previous simulations with a similar model demonstrating why a phonological impairment would bring about impaired development of the knowledge of spelling-sound correspondences. Then we consider why remediations targeting phonological skills lead to substantially reduced benefit in older children, and finally why remediations that explicitly target such correspondences have greater success.

This framework allows us to examine the mapping hypothesis in a new way. The hypothesis holds that phonological awareness interventions influence word recognition processes via the quality of the letter-sound mapping representations (rather than by the quality of the phonological representations per se). This was illustrated computationally in the Harm and Seidenberg (1999) paper by showing that the internal representations of similar words (MEAT, TREAT, HEAT) were further apart in a phonologically impaired model, even after more experience accrues. Further analysis of the Harm and Seidenberg (1999) model demonstrated that the phonologically impaired model was

intervention techniques were employed. Furthermore, we hypothesized that the Word Building technique might place a strong computational pressure on the system to systematically represent the contribution of each grapheme to the pronunciation of the words.

more sensitive to the M in MEAT than the normal model (see also Siegel, Share, & Geva, 1995, for evidence of dyslexics attending to word specific aspects of print more so than controls). The mapping hypothesis leads to several predictions that can be tested by exploring the model in an intervention framework.

When applied as a pre-reading preventative measure, phonological awareness training should help children (or models) with poor phonological representations to form letter-sound mapping representations that are more highly componential. Thus, repairing phonology before the beginning of reading instruction should lead to the representations of similar items such as MEAT, TREAT, HEAT overlapping more, because it alleviates the pressure to learn disparate representations for such items.

However, when applied as a late remediation technique, phonological awareness should help children remediate their phonological representations, but should put little pressure on the system to change the letter-sound mappings that are already established. Thus repairing phonology after reading starts should not lead to improvements in non-word generalization. The explanation provided by the mapping hypothesis is that letter-sound mappings will remain unchanged, and there will be no pressure for MEAT, TREAT, HEAT to move closer together. Finally, interventions like Word Building (or the class of interventions that places emphasis on letter-sound units) should improve non-word generalization by making spelling to sound mappings more componential.

A Connectionist Model of Reading Development

Here we provide an overview of the model and core results from Harm and Seidenberg (1999). As noted earlier, this model differed from previous simulations of reading (e.g., Seidenberg & McClelland, 1989; Plaut et al., 1996; Coltheart, Curtis, Atkins, & Haller, 1993; Coltheart et al., 2001) in that it incorporated a trainable phonological system as the output of the model. This phonological system was implemented as a set of low level phonetic features such as voicing, plosive, etc. derived from linguistic theory, with a set of weighted connections between such features. Additionally, a set of “cleanup” units received input from all phonological units, and fed activations back to all such units; the “cleanup” units allowed the phonological system to learn higher order relationships between phonemic features. Figure 1 depicts the phonological system.

The phonological system was trained prior to addition of the reading component. This training corresponds to the knowledge of the sound structure of language that children have prior to literacy instruction. Details of this training procedure are provided in Harm and Seidenberg (1999); for our purposes here it is sufficient to say that the model was trained to learn the relationships between phonolog-




Figure 1. The phonological system used in Harm and Seidenberg (1999).

ical units in the target language. This included the inventory of possible phonemes, and phonotactic constraints between such phonemes in sequence. For example, the model learned, through training, that in English, two nasal phonemes cannot precede a vowel (e.g., /mnoʊ/).

Several variants of the phonological system were developed; a normal system which was unhindered in its ability to learn such relationships, and various impaired versions which, through manipulation of the systems architecture or learning regime, were limited in their capacity to learn such phonological constraints. Mild, moderate and severe impairments were all explored, and the level of impairment was related to performance on tasks such as phoneme restoration and categorical perception of phonemes.

Having created a set of phonological systems with varying knowledge of the sound structure of English, Harm and Seidenberg (1999) then created a set of reading models in which these pre-structured phonological systems were used as the starting state for the reading task. The architecture of the entire reading model is shown in Figure 2.

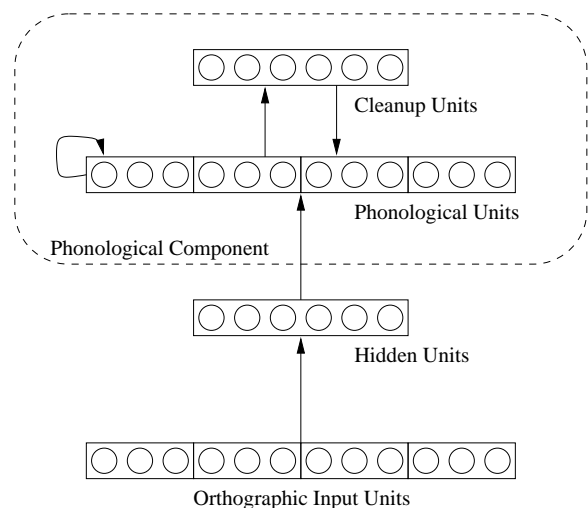


Figure 2. The intact reading model from Harm and Seidenberg (1999).

The phonologically impaired models all demonstrated impairments in nonword reading, with nonword reading performance being inversely correlated with the degree of

phonological impairment. Varying the degree of phonological impairment provided a close fit, in word and nonword reading skill, to a wide range of subjects in the Manis et al. (1996) study.

Close examination of the model's internal learned representations gave insights into why a phonological impairment leads immediately to nonword reading deficits, and ultimately to word reading deficits as well. The normal, unimpaired model's phonological system was able to repair partial or noisy results that were produced by the reading system; this lessened the workload imposed upon the hidden units between the orthographic representation and phonological system. If the hidden units produced an output that was not exactly correct, but was sufficiently "in the ballpark" then the phonological system could pull that partial result into a correct representation. For the simulations with an impaired phonological system, more work had to be done by these hidden units, because the phonological system was less able to repair partial results. This increased workload on the hidden units caused them to be more likely to "memorize" word forms, and form item-specific representations.

As one concrete example, Harm and Seidenberg (1999) considered a cluster of words ending in -EAT (e.g., MEAT, SEAT, EAT, TREAT) plus a nonword GEAT. It was found that the hidden unit representations for these words overlapped with each other less in the impaired simulations than in the normal one (Harm & Seidenberg, 1999, Figures 15-18). For example, the impaired simulations showed a much stronger influence of the letter M on the vowel pronunciation of MEAT for the impaired simulations than the normal ones (Harm & Seidenberg, 1999, Figure 19). The pronunciation of the vowel in MEAT should not require attending to the letter M at the beginning, as its pronunciation is predictable from the orthographic rime. Because the impaired simulations produced more item-specific representations of EAT, MEAT, SEAT and TREAT, it was not able to correctly pronounce the nonword GEAT; this nonword did not fall within the cluster of words with the same rime. For the normal simulation, GEAT did overlap considerably with other -EAT words, and hence that simulation could pronounce that nonword.

The model thus provided a computational account of why poor phonological representations lead to poor reading, and in particular poor nonword generalization. The crucial insight from these simulations is that a phonological impairment leads to poor learning in the orthography→phonology component. Instead of forming representations sensitive to subword units such as onsets and rimes, the hidden units in the impaired simulations learn item-specific representations. The formation of these item-specific representations is what directly impairs nonword reading. The poor nonword reading in the model is not due to the phonological system's impaired ability to assemble phonemes produced by the reading system, but

rather the phonological impairment causes poor orth→phon representations to be formed during learning.

This makes an important prediction: once the poor representations are formed during learning, repairing the phonological system will not change these representations. Additionally, learning in models such as this one is parasitic on existing representations, so once such poor representations in the hidden unit layer are in place, they will influence the types of representations formed when subsequent words are acquired. Consistent with observed empirical studies, this analysis predicts that remediations attempting to improve phonological skills will not lead to improvements in nonword reading, if such interventions take place after poor, item-specific representations have become solidified in the hidden unit layer. For an intervention to be successful, it would have to break such item-specific representations and force the formation of more componential ones. This analysis provides an answer to the puzzle posed earlier: poor phonological representations cause poor representations to be formed in the orth→phon system, yet repairing the phonological representations alone will not repair the orth→phon system once such representations have become entrenched.

We conducted a series of simulations to explore these predictions. First we consider interventions that seek to repair phonological representations, and then turn to ones designed to promote more componential representations in the mapping from spelling to sound.

Simulation 1: Remediating Phonology

In this simulation, we began with an impaired phonological representation identical to that used in Harm and Seidenberg (1999). The phonological impairment was alleviated at different points in reading instruction, to determine not only if this led to improved reading, but how sensitive this improvement was to the time at which the intervention was applied.

Method

Items. The same training corpus, phonological representation and orthographic representation were used as that in Harm and Seidenberg (1999). This consisted of 3,123 uninflected monosyllables with six phonological slots, one slot for each phoneme. Each slot consisted of eleven phonemic features, for a total of 66 phoneme feature units. A set of 208 orthographic units were used, consisting of 26 units (corresponding to the letters A-Z) for each orthographic slot, and a total of eight orthographic slots.

Procedure. The phonological impairment involved the lesioning of a random 50% of the connections within the phonological system (in practice, setting the weight value to zero and freezing the connections to prevent them from

changing). Further, weight decay was imposed on the remaining connections, such that upon presentation of each word, the weights all decay towards zero at a rate proportional to 0.00001 multiplied by the weight's current value. This limits the maximum magnitude which the weights can attain.

To simulate a reparation of the phonological representations, these impairments were ceased during reading training. Specifically, weight decay was discontinued, and the previously lesioned connections were initialized to small random values and unfrozen (as they were at the start of training for the normal model), so they could begin developing. The goal of this simulation was not to test a particular method of remediating phonology, but rather to ask, in the limit, how much of a benefit in reading performance can be gained if one could (by whatever means) totally alleviate a phonological impairment.

We created a total of five simulations: a normal simulation with intact phonological representations, an impaired simulation, and three remediated simulations which began training with a phonological impairment which was then alleviated as described above at different points in development. For one of these remediated simulations we alleviated the impairment to the phonological representations at the beginning of reading instruction, for another after 10,000 reading trials (the presentation of a single word counting as a trial), and the final one after 100,000 word presentations. At the very beginning of training the normal model could not accurately read any words in the training set; at 10,000 word presentations it could correctly read 12% of the training set items; at 100,000 it could read 41%. These three points in training were chosen for the intervention because they represent a reasonable range of early reading competence.

These conditions allowed us to explore the impact of remediating phonological impairments at the onset of literacy training, early in training, and later in training, to determine the impact on when interventions began their effectiveness. As in Harm and Seidenberg (1999), all models were trained for a total of 10 million word presentations.

Scoring. Materials from the Woodcock Word Attack Form H (Woodcock 1987; a standardized test of nonword reading) and the Woodcock Word Identification (Form H) were used to assess the models' word and nonword reading. The 35 monosyllabic nonwords from the Word Attack and 39 monosyllabic words from the Word Identification test were used. For each phonological slot, the model's output was compared with the representation of all existing phonemes, and the phoneme closest to the model's output was considered the model's output. The output was considered correct if all phonemes matched the target, and incorrect if any of the phonemes did not match.

Results

The number of correct items from the Woodcock Word Attack (using 35 items; polysyllabic items were excluded) are shown in Figure 3. When the phonological impairment was alleviated at the beginning of reading instruction, the model's performance was nearly identical to that of the normal simulation (the difference is just one item for Woodcock Word attack). However at later points in training (even as early as 10,000 word presentations), the intervention did not improve nonword reading; the accuracy was similar to the impaired, unremediated model.

For word reading accuracy, all models scored 37-38 correct of the 39 monosyllabic items in the Woodcock Word Identification test. The differences in performance were most clear in nonword reading.




Figure 3. Results of Simulation 1. Alleviating phonological impairment at the onset of literacy training (OK iterations) resulted in normal performance; any point later in training resulted in impaired nonword reading.

Discussion

This simulation demonstrates that in the model, interventions that target phonological representations have a potential for success provided they are introduced extremely early in learning. When phonology was repaired at 10K iterations it is already too late for this intervention to ultimately produce normal nonword reading skill. This is because in models such as this, subsequent learning is a function of what has already been learned; repairing phonological representations once poor learning has started to become entrenched has a much lower chance of success.

Simulation 2: Simulating the Word Building Intervention

McCandliss et al. (in press) examined the reading skills of children, age 7-10, who had deficient decoding skills, and traced their progress across 20 sessions of a decoding-skills intervention called Word Building, an intervention specifically designed to place a strong emphasis on systematic letter-sound relationships, and minimal direct focus on analysis of spoken words. Initially, children in this study exhibited deficits in decoding and phonemic awareness skills, suggesting that these children were not engaging in full alphabetic decoding. The intervention focused attention on each grapheme position within a word through a procedure of progressive minimal pairing of words that differed by one grapheme (see Figure 4). Children worked on a given word until they correctly pronounced it; then they moved on to the next word in a lesson. If a word was mispronounced, the child's attention was directed to the components of the word to emphasize the visual overlap with preceding items. A new lesson was begun once the child successfully pronounced all items in the current lesson. Relative to a randomized control group, children assigned to the intervention demonstrated significantly greater improvements in standardized measures of decoding and phonological awareness.

Given the success of the McCandliss et al. (in press) intervention, the purpose of this simulation was first to determine whether the simulation of this intervention was similarly successful, and second to determine whether, as the mapping hypothesis predicts, the reason the intervention succeeded was because it induced changes in the componential nature of spelling to sound mappings.




Figure 4. Sample stimuli from the McCandliss et al. (in press) Word Building intervention. Consecutive items in the sequence were created by changing or moving only one grapheme.

Method

Items. The items and representations used in this simulation were identical to those used in Simulation 1.

Procedure. In these simulations, we introduced the Word Building remediation at the onset of literacy training, after 10,000 word presentations, and after 100,000 word presentations, as in Simulation 1. The intervention worked as follows. Words were presented to the model as normal, except that for every word presentation, there was a 0.1% chance that a “lesson” would be taken.² The lessons consisted of the exact words in the same order as were used in the McCandliss et al. (in press) Word Building intervention. All words in the lesson were also present in the main training corpus, so no additional unique items were supplied by the lessons. The model began at lesson 1, and kept track of the current lesson. As in the intervention, items in a lesson were processed in order, with the error recorded. A lesson was considered complete if every word generated an overall sum squared error of less than 1.0; once a lesson was complete, the model recorded which lesson it was to use next, and resumed normal training, to begin the next lesson once in every 1,000 normal word training trials. As in the intervention with children, if there was an error on a given word in a lesson, then the word was broken into constituent units. In the model, this was implemented by breaking the word into onset and rime units. The model received a training trial on the onset in isolation, and then the rime in isolation. For example, if the model generated an error > 1.0 for CAT, it would receive a training trial for C in the corresponding letter position mapping to /k/ in the phonological representation. Then it would receive a training trial for AT mapping onto /æt/ in the phonological representation. When all lessons had been completed the model began again at lesson 1.³ There were 77 distinct lessons, covering a total of 445 unique monosyllabic words.

Accuracy was tested at the completion of training, and assessed in the same manner as Simulation 1.

Results

Figure 5 shows the Woodcock Word Attack accuracy for the normal and impaired models, and the three remediated models, where remediation began at the beginning of literacy training (zero K iterations), 10,000 (10 K) iterations and 100,000 (100 K) iterations.

²In pilot studies we tried a variety of probabilities for a lesson to be encountered, ranging from 0.05% to 0.3%, and found that it did not have a qualitative affect on the results. A training probability of 0.1% is small enough that the extra presentation of words from the intervention is very unlikely to have any qualitative effect on the normal distribution of words.

³The model went through more sweeps of the lessons than the children in the McCandliss et al. (in press) did, however subsequent passes through the lessons were largely reminder trials. It was only on the first pass through the set of lessons that the model generated errors that required breaking apart the words for componential training. All models completed the first sweep of the set of lessons within 250,000 word presentations.

In contrast to Simulation 1, where the intervention only was successful when applied at the beginning of literacy training, here the intervention was successful at the onset of literacy training, early in training and later in training. When the intervention began at 10K and 100K the improvement in reading level over the impaired simulation was 0.8 grade levels; very close to the magnitude of improvement reported by McCandliss et al. (in press).

We also examined the degree to which intervention-based improvements might be driven by learning specific rime patterns that appear within the intervention. Of the 35 monosyllabic items in the Woodcock Word Attack Form H, only 15 contained phonological rimes that appeared in the lessons (e.g., the nonword PHET, which shares the same rime as lesson items SET, MET, PET). Next we examined just the test items that discriminated the impaired-remediated model from the impaired non-remediated model. Of the items that discriminated the two models at 10K, only 40% contained rimes that appeared in the lesson. Of the items that discriminated the two models at 100K, only 18% of the items contained rimes that appeared in the lessons. This suggests that the effect of the lessons extended beyond the actual rimes taught in the lessons, and into more general properties of decoding words. We now explore the effect of the training on the models' internal representations in greater detail.




Figure 5. Results of Simulation 2. Imposition of remediation regime had a beneficial effect even late in training.

Discussion and Analysis

In the Harm and Seidenberg (1999) model, the effect of poor phonological representations was to cause the hidden units between orthography and phonology to develop

representations that are less componential and more holistic. Specifically, the variability between representations of words with similar orthographic and phonological rimes (e.g., CAT, BAT, MAT) was much greater in the impaired simulation. The normal model had much more similar representations of such words.

In analyzing the hidden unit representations, we observed that for many words sharing common spellings, the representations tended to overlap more in the normal and remediated models than in the impaired one. To illustrate this effect, we first created multidimensional scaling plots of a set of words and one nonword with a common orthographic rime (EAT, MEAT, SEAT, FEAT, HEAT, BEAT, GREAT, THREAT and the nonword GEAT). The multidimensional scaling algorithm takes relative distances between high dimensional sets of numbers (in our case, 100 numbers representing the activity of the 100 hidden units) and represents those relative distances on an arbitrary two dimensional scale. Using this technique, the relative overlap in hidden unit activities for these words was plotted for the normal, impaired and remediated model.

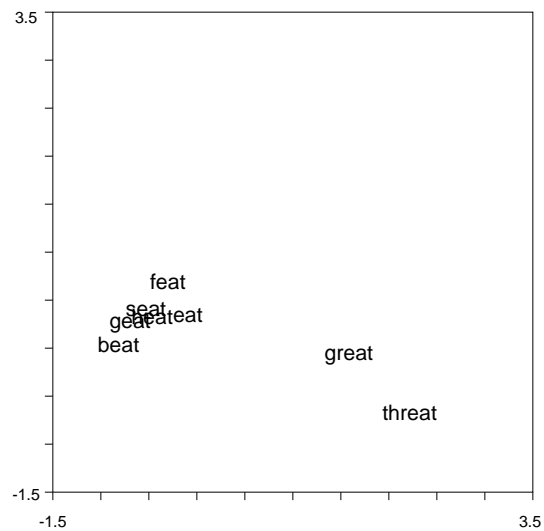


Figure 6. Multidimensional scaling plot of normal model. Representations of words with similar spellings and pronunciations (EAT, BEAT, MEAT, SEAT, HEAT) densely clustered. See text.

Figure 6 shows the plot for the normal model. The words and nonword with the same pronunciation are tightly clustered, while the exception words GREAT and THREAT are located further away in representation space. The impaired model's plot shown in Figure 7 shows a much more even spreading of words throughout representational space; the words with the same orthographic and phonological rime are not very much closer to each other than they are to the exception words with a different phonological representation. For example, in the impaired simulation (Figure 7)

the word HEAT is about as far from the related word BEAT as it is from GREAT. Although the normal model could pronounce the nonword GEAT, the impaired one could not, because the pattern of hidden unit activity generated by GEAT was not close enough to the representation of similar words. The impaired model could pronounce the word HEAT, but had formed sufficiently item-specific representations of it that the nonword GEAT could not take advantage of any overlap.

Figure 8 shows the same plot for the model remediated at 10,000 word presentations.⁴ This data qualitatively matches that of the normal model, in that the similarly spelled regular words (and the nonword) were more densely clustered; while the exception words are set further apart from the regular words. Here, as in the normal model, the nonword GEAT could be successfully pronounced.

The tighter clustering of words in the normal and remediated model also makes word reading easier. Harm and Seidenberg (1999) found that with mild phonological impairments, word reading was still at relatively normal levels (as was found with the few children fitting this description in the study by Manis et al., 1996). However, with more extreme phonological impairments, additional resources must be recruited to “memorize” each individual word, and this resulted in poor performance on the most difficult words, the low frequency exceptions. These simulations focused on the more moderately impaired simulations; we discuss issues of word identification in greater detail in the General Discussion.

The multidimensional scaling plots are useful for visualizing the changes in the representations as a result of the remediation. However, it is not clear from such demonstrations how robust such changes are across larger sets of items. To test for this phenomenon more broadly, we algorithmically created a list of all pairs of words in the training corpus having the same orthographic rime and the same phonological rime. A set of 12,416 such pairs resulted. As a control, a set of 12,416 random word pairs were also generated, with the only constraint being that members of each pair did not share the same orthographic or phonological rime. The final state of the impaired model was run on all word pairs, with the activity in the hidden unit layer being recorded. This procedure was then repeated for the final state of the model where remediation began at 10,000 word presentations.

The euclidean distance (a measure of distance between high dimensional points in space) between the hidden unit activation for each member of a pair was recorded, giving 12,416 differences for the rhyming words, and the same number of differences for the random non-rhyming pairs. If the effect of the intervention was to pull similar representa-




Figure 7. Multidimensional scaling plot of the impaired model. Words with similar spellings and pronunciations were not well clustered relative to their distance from words with similar spellings but different pronunciations (GREAT, THREAT).




Figure 8. Multidimensional scaling plot of the model remediated with Word Building technique. As in Figure 6, the words with similar spellings and pronunciations closely overlapped in representational space.

⁴The models remediated at other points in training showed qualitatively similar results, so to conserve space we only plotted the data from the model remediated at 10,000 presentations.

tions closer and disparate representations further apart, then a two way interaction should be observed between list type (rhyming or random) and presence of intervention.

Figure 9 shows the mean difference in closeness of representation following intervention for the rhyming words and the random words. The difference for the rhyming words was reliably below zero ($t(12416) = -14.9, p < 0.001$), indicating that the intervention had decreased the hidden unit distance between rhyming words. Conversely, the change in hidden unit space for the random items was reliably above zero ($t(12416) = 18.46, p < 0.001$), indicating that the remediation had pulled dissimilar items apart.




Figure 9. Changes in hidden unit differences as a result of the intervention, for rhyming words and random word pairs. Rhyming words were closer together as a result of the intervention, while random unrelated word pairs moved further apart.

Among the rhyming words, there is variability in the extent to which the intervention pulled the representations closer to each other. The Word Building lessons contained a number of words representing certain rime classes, such as -AT. Other rime classes, such as -OUGH in ROUGH and TOUGH were not present in the lessons. We further analyzed the rhyming words, breaking them apart into two classes: those pairs of words containing rimes found in the lessons (such as CAT/BAT), and those containing rimes not in the lessons (such as ROUGH/TOUGH). The purpose of this analysis was to determine both if there was a reliable difference between the rime units in the lessons and those not, and to see if words not present in the lessons benefited at all from the intervention.

Figure 10 shows the results. Not surprisingly, the trained items (those from rime units found in the lessons) showed a reliable shift towards reducing the differences in hidden unit representations ($t(5061) = -14.46, p < 0.001$).

Interestingly, the rime pairs containing rimes not covered by the lessons also showed a reliable, though smaller difference ($t(7355) = -7.43, p < 0.001$), indicating that even though they were not represented explicitly by the lessons their representations still moved closer together as a result of the intervention. Such results suggest that interventions of this form are most effective for rime units explicitly contained within the lessons, but that there is a crossover effect to other rime classes.




Figure 10. Changes in hidden unit differences as a result of intervention, for rhyming words containing a rime that was part of a lesson in the intervention, and rhyming words containing a rime that was not found in any lesson in the intervention.

General Discussion

In this paper we have provided a mechanistic account of two effects in the literature that, until now, have been reported as empirical observations without a theory-based explanation (i.e., based on tenets of the phonological deficit hypothesis.) The first is that the effectiveness of interventions that target phonological representations through speech activities declines sharply once exposure to print has begun. In particular, we have shown that the effect of poor phonological representations begins to influence the development of the orth→phon system very early on in training. Even if a purely phonological intervention can bring phonological skills up to a normal level, the effectiveness of beginning such an intervention declines sharply once development of poor representations in the orth→phon system has begun. After only 10,000 word presentations, it is already too late to undo the effect of learning with a poor phonological representation. The second effect we have simulated is that interventions targeting the relationship between print and sound produce improvements in nonword reading when

applied at either early or later stages in development. Such improvements result from increasing the representational overlap of words with common rimes, and decreasing the overlap between arbitrary words. Additionally, the simulations show that there can be some transfer of learning the rimes contained in the lesson to other rimes.

We have presented the mapping hypothesis as a means of accounting for these puzzling results concerning the relative effectiveness of different intervention techniques. This work leverages the findings from a recent set of simulations that explicitly model the role of phonological skills in reading development (Harm & Seidenberg, 1999). It represents the first step in what promises to be an exciting new branch of research in developmental reading impairments: simulating detailed aspects of the child's phonological knowledge and experience with print, examining changes that occur in the structure of representations over time, and making use of this explicit computational model to provide systematic and novel accounts of empirical findings in the intervention literature.

The results of Simulation 1 demonstrate that phonological impairments, if addressed early enough, can bring about improvements in nonword reading. Simulation 2 demonstrates that interventions that target the componential aspects of words also produce improvements in nonword reading. Future work can examine how these factors interact, and the extent to which combinations of the two produce additional benefits (as was found in Ball & Blachman, 1991). Additionally, future simulations can examine whether the magnitude of phonological impairment interacts in important ways with the nature and timing of interventions.

In the current work, we used corpora of monosyllables weighted by their frequency of occurrence in printed text. In other work (Foorman et al., 2001; Harm & Seidenberg, 2002) we have made preliminary explorations of the effect of different reading basals by using actual texts. In future work we hope to combine the techniques of simulating interventions with the use of materials throughout training that more closely match the child's experience.

When considering the implications of this work for understanding the impact of reading interventions, one consideration that comes to the forefront involves understanding intervention-based changes in word identification performance. As a first step, this paper has focused on interventions that do and do not yield improvements in nonword reading, specifically to focus investigation on the emergence of systematic representations that might generalize to novel words. Data concerning transfer of intervention benefits from short-term improvements in non-word reading to long-term improvements in word reading assays are mixed (see Torgesen et al., 2001 and Wise, Ring, & Olson, 2000; but see also National Reading Panel, 2000; see Olson, in press for discussion). An important step in future work along these lines will be to investigate the complexities en-

countered when measuring improvements through assays of word reading abilities. Developmental/intervention modeling efforts might provide an additional avenue of investigation into the complexities of this form of transfer.

There are, however, several complexities that arise in the use of Word Identification tests as a measure of improvements for decoding intervention programs. For example, many standard tests of word reading, such as the WRMT-R and Woodcock-Johnson Word Identification subscale contain a high proportion of "exception" words that might benefit less from incremental improvements in decoding abilities (see McCandliss et al., in press, for discussion). One potential advantage of the modeling approach that could be exploited in future studies is that models never tire of testing, thus enabling researchers to test improvements in word reading abilities on a large corpus of words that are representative of text experience, rather than on standardized tests that might select a small number of items that maximize the tests' sensitivity to individual differences at the expense of providing materials that accurately represent reading materials.

Secondly, tests of word identification are often more steeply graduated for monosyllabic words than for their pseudoword counterparts, and may be less sensitive as a measure of improvements in monosyllabic word reading.⁵ These complexities prevented us from using such measures productively in this current investigation, which used a model designed to account for data on monosyllabic words. Future work might use a more exhaustive test of real words to assess change, and might also expand the model to process multi-syllabic words.

Another complexity introduced when considering abilities in identifying previously encountered words is highlighted when contrasting the current model to the "triangle" model of word reading (Seidenberg & McClelland, 1989, see Figure 11). In the triangle model, there are two paths from print to pronunciation; orth→phon and orth→sem→phon. Words have the potential to utilize both pathways to a much greater extent than nonwords do; nonwords, broadly, can only use the orth→phon pathway (which the present simulations have focused on). Analyses of the "triangle" model (Harm & Seidenberg, submitted) revealed that the orth→sem pathway is far less tuned to the componential nature of print, and is more "holistic" in its processing of words, because, with the exception of morphemes, correlations between the components of print and meaning are generally not very systematic (e.g., the B in BAT gives a clue to the initial pronunciation of the word, but gives very little hint as to its meaning).

Thus there are two ways in which a model with poor

⁵For example, of the 106 items in the Woodcock Word Identification Form H, only 43 are monosyllabic, with 31 of those items appearing in the first 35 (least difficult) items. A score of 35 correct corresponds to a reading grade equivalence of 2.0.




Figure 11. The Seidenberg and McClelland (1989) model of visual word recognition.

nonword reading could be described as relying on “global” processing strategies: the formation of holistic, word specific representations in orth→phon, and an over-reliance on the orth→sem→phon pathway, which by its nature is more holistic and tuned to global processing.

One prediction of this theoretical framework is that interventions that target whole word reading, particularly embedded in a semantic context, could yield improvements in word reading equal to those that focus on sublexical units. Consistent with this, a study by Olson and colleagues (Olson, Wise, Ring, & Johnson, 1997; Wise et al., 2000) found that interventions that focused on sublexical units yielded reliably better nonword performance than those with an equal amount of print exposure, but focusing on whole words; however there was no advantage in word reading for one intervention over the other. This result could come about because in the latter (whole word) intervention, the orth→sem→phon pathway was being exercised, while in the former, the orth→phon pathway experienced the benefit (resulting in better nonword reading as well).

Further, studies suggesting that poor readers form “word specific” representations that do not generalize well (e.g., Stanovich & Siegel, 1994; Byrne, 1992) can be explored within this framework, potentially offering diagnostic tests to determine the exact locus of “word specific” representations, and suggesting (potentially different) interventions tuned to the specific source of the impairment.

Conclusion

Our models are consistent with behavioral evidence indicating that the nature of phonological representations is shaped by knowledge of orthography (e.g., Morais et al., 1979). Knowledge of phonemes in particular seems intimately tied to knowledge of orthography, and our models

shed some light on why this is so. In such models phonological representations are shaped by their participation in different tasks (reading, listening, speaking; Plaut & Kello, 1999; Harm & Seidenberg, 1999). If knowledge of phonemic structure is critical to skilled reading, and this knowledge is normally acquired mainly through the pairing of orthography and phonology, then teaching methods and remedial interventions that emphasize this pairing should be effective. Computational simulations such as those presented here can be a valuable tool for integrating existing empirical findings within an explanatory framework—one that can be productively explored to inform further empirical research.

References

- Adams, M. J., Treiman, R., & Pressley, M. (1998). Reading, writing and literacy. In I. E. Sigel & K. A. Renninger (Eds.), *Handbook of child psychology. vol 4: Child psychology in practice* (p. 275-355). New York: Wiley.
- Ball, E. W., & Blachman, B. A. (1991). Does phoneme segmentation training in kindergarten make a difference in early word recognition and developmental spelling? *Reading Research Quarterly, 24*, 49-66.
- Beck, I. (1989). *Reading today and tomorrow: Teachers edition for grades 1 and 2*. Austin, TX: Holt and Co.
- Blachman, B. A., Ball, E. W., Black, R. S., & Tangel, D. M. (1994). Kindergarten teachers develop phoneme awareness in low-income, inner-city classrooms. Does it make a difference? *Reading and Writing, 6*, 1-18.
- Bradley, L., & Bryant, P. (1983). Categorizing sounds and learning to read - A causal connection. *Nature, 301*, 419-421.
- Bus, A. G., & Ijzendoorn, M. H. van. (1999). Phonological awareness and early reading: A meta-analysis of experimental training studies. *Journal of Educational Psychology, 91*, 403-414.
- Byrne, B. (1992). Acquisition of literacy. *International Journal Of Psychology, 27*, 565-565.
- Byrne, B., & Fielding, B. (1995). Evaluation of a program to teach phonemic awareness to young children - a 2-year and 3-year follow-up and a new preschool trial. *Journal Of Educational Psychology, 87*, 488-503.
- 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., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). Drc: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review, 108*(1), 204-256.
- Cunningham, A. E. (1990). Explicit versus implicit instruction in phonemic awareness. *Journal of Experimental Child Psychology, 50*, 429-444.
- Ehri, L., & Nunes, S. (in press). *Phonemic awareness instruction helps children learn to read: Evidence from the national reading panel's meta-analysis*. (in press, *Reading Research Quarterly*)
- Foorman, B. R., Perfetti, C., Seidenberg, M., Francis, D., & Harm, M. (2001). *What kind of text is a decodable text? And what*

- kind of text is an authentic text?* (Talk presented at the Annual Meeting of the American Education Research Association, Seattle, WA, April 2001)
- Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*(3), 491-528.
- Harm, M. W., & Seidenberg, M. S. (2002). *Modeling reading development from first grade text*. (Talk presented at the Annual Meeting of the American Education Research Association, New Orleans, LA, April 2002)
- Harm, M. W., & Seidenberg, M. S. (submitted). *Computing the meanings of words in reading: Division of labor between visual and phonological processes*. (Manuscript submitted for publication)
- Hatcher, P. J., Hulme, C., & Ellis, E. W. (1994). Ameliorating early reading failure by integrating the teaching of reading and phonological skills: The phonological linkage hypothesis. *Child Development*, *65*, 41-55.
- Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, *98*(1), 74-95.
- Lieberman, I., Shankweiler, D., Fisher, F., & Carter, B. (1974). Explicit syllable and phoneme segmentation in the young child. *Journal of Experimental Child Psychology*, *18*, 201-212.
- Lundberg, I. (1994). Reading difficulties can be predicted and prevented: A Scandinavian perspective on phonological awareness and reading. In C. Hulme & M. Snowling (Eds.), *Reading development and dyslexia* (p. 180-199). London: Whurr.
- Lundberg, I., Frost, J., & Petersen, O. (1988). Effects of an extensive program for stimulating phonological awareness in preschool children. *Reading Research Quarterly*, *23*, 263-284.
- Lundberg, I., Olofsson, A., & Wall, S. (1980). Reading and spelling skills in the first school years predicted from phonemic awareness skills in kindergarten. *Scandinavian Journal of Psychology*, *21*, 159-173.
- 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.
- Mann, V. A. (1984). Longitudinal prediction and prevention of early reading difficulty. *Annals of dyslexia*, *34*, 115-136.
- McCandliss, B. D., Beck, I., Sandak, R., & Perfetti, C. (in press). *Focusing attention on decoding for children with poor reading skills: A study of the Word Building intervention*. (To appear in *Scientific Studies of Reading*)
- Morais, J., Alegria, J., & Content, A. (1987). The relationship between segmental analysis and alphabetic literacy: An interactive view. *Cahiers de Psychologie Cognitive*, *7*(5), 415-438.
- Morais, J., Cary, L., Alegria, J., & Bertelson, P. (1979). Does awareness of speech as a sequence of phones arise spontaneously? *Cognition*, *7*, 323-331.
- National Reading Panel. (2000). *Teaching children to read: An evidence-based assessment of the scientific research literature on reading and its implications for reading instruction*. (National Institute of Child Health and Human Development, Washington, D.C.)
- Olson, R. (in press). *Phoneme awareness and reading, from the old to the new millenium*. (To appear in E. Hjelmquist (Ed.), *Dyslexia and literacy: A tribute to Ingvar Lundberg*)
- Olson, R., Wise, B., Ring, J., & Johnson, M. (1997). Computer-based remedial training in phoneme awareness in phonological decoding: Effects on the post-training development on word recognition. *Scientific Studies of Reading*, *1*, 235-253.
- Plaut, D. C., & Kello, C. T. (1999). The interplay of speech comprehension and production in phonological development: A forward modeling approach. In B. MacWhinney (Ed.), *The emergence of language* (p. 381-415). Mahweh, NJ: Erlbaum.
- 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.
- Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, *10*(5), 377-500.
- Read, C., Yun-Fei, Z., Hong-Yin, N., & Bao-Qing, D. (1987). The ability to manipulate speech sounds depends on knowing alphabetic writing. In P. Bertelson (Ed.), *The onset of literacy* (p. 31-44). Cambridge, MA: MIT Press.
- Rosner, J., & Simon, D. P. (1971). The auditory analysis test: An initial report. *Journal of Learning Disabilities*, *4*(7), 40-48.
- Schneider, W., Ennemoser, M., Roth, E., & Kaspert, P. (1999). Kindergarten prevention of dyslexia: Does training in phonological awareness work for everybody? *Journal of Learning Disabilities*, *32*, 429-436.
- Schneider, W., Roth, E., & Ennemoser, M. (2000). Training phonological skills and letter knowledge in children at risk for dyslexia: A comparison of three kindergarten intervention programs. *Journal of Educational Psychology*, *92*(2), 284-295.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523-568.
- Siegel, L. S., Share, D., & Geva, E. (1995). Evidence for superior orthographic skills in dyslexics. *Psychological Science*, *6*(4), 250-254.
- Stanovich, K. E., & Siegel, L. S. (1994). Phenotypic performance profile of children with reading disabilities - a regression-based test of the phonological-core variable-difference model. *Journal Of Educational Psychology*, *86*, 24-53.
- Torgesen, J. K., Alexander, A. W., Wagner, R. K., Rashotte, C. A., Voeller, K., & Conway, T. (2001). Intensive remedial instruction for children with severe reading disabilities: Immediate and long-term outcomes from two instructional approaches. *Journal Of Learning Disabilities*, *34*, 33-58.
- Tunmer, W. E., & Nesdale, A. R. (1985). Phonemic segmentation skill and beginning reading. *Journal of Educational Psychology*, *77*, 417-427.
- Wagner, R. K., & Torgesen, J. K. (1987). The nature of phonological processing and its causal role in the acquisition of reading skills. *Psychological Bulletin*, *101*, 192-212.
- Wagner, R. K., Torgesen, J. K., & Rashotte, C. A. (1994). Development of reading-related phonological processing abilities - new evidence of bidirectional causality from a latent variable longitudinal-study. *Developmental Psychology*, *30*, 73-87.
- Weiner, S. (1994). Effects of phonemic awareness training on low-achieving and middle-achieving 1st-grader's phoneme awareness and reading ability. *Journal of Reading Behavior*, *26*(3), 277-300.

- Wise, B., Ring, J., & Olson, R. (2000). Individual differences in gains from computer-assisted remedial reading. *Journal Of Experimental Child Psychology*, 77, 197-235.
- Woodcock reading mastery tests - Revised*. (1987). (Circle Pines, MN; American Guidance Service)