A Connectionist Model of Alphabetic Spelling Development and Developmental and Acquired Dysgraphia

Richard P. W. Loosemore, Gordon D. A. Brown and Frances L. Watson

Cognitive Neurocomputation Unit
Department of Psychology
University of Wales, Bangor
Gwynedd, LL57 2DG
United Kingdom
Earn/BITNET: PSS030@uk.ac.bangor.vaxa

Abstract

In this paper we describe a connectionist model of the development of alphabetic spelling. The model learns to spell regular words more quickly than words with irregular spellings. When the computational resources available to the model are restricted, the model learns more slowly and, analogously to developmental dyslexics, fails to learn some of the irregular items in its vocabulary. Experimental evidence is reported, which shows that both normal and dyslexic children of various ages have difficulty with particular word types that are similar to the problems experienced by the model on the same words. Finally, the model is "lesioned," and its performance is then similar to that of "surface dysgraphics." The good fit between model and data is taken as evidence that, throughout much of the relevant developmental period, the task facing children can be usefully viewed as a statistical one.

Introduction

Our purpose in the present paper is to apply the connectionist modelling approach, which has provided useful unifying models of spelling-to-sound translation in reading (e.g. Seidenberg & McClelland, 1989a), to the case of spelling development. Here, the computational task faced by children is one of mapping phonology onto orthography. This mapping is more "irregular" than the mapping from spelling to sound investigated by Seidenberg & McClelland, and one of our aims is to determine whether a connectionist system could learn this more demanding set of associations.

The spelling of many English words cannot reliably be derived from phonological information alone, because of words such as SOAP (cf. ROPE, COPE, HOPE), and pairs or triples of words that are pronounced identically but spelled differently (e.g. BARE/BEAR and HAIR/HARE).

Pronunciation information nevertheless partially constrains the written forms of words, and there is abundant psychological evidence, from the study of adult and child spelling errors, that pronunciation information is used in spelling. The spelling of non-words is also generally assumed to implicate mechanisms for translating sound into spelling, as are the characteristic patterns of breakdown exhibited by patients suffering from various forms of head injury (see Shallice, 1988, for a review).

These considerations have lead psychologists to produce models of spelling in which there are two separate routines available to translate sound into spelling. Early spelling is assumed to occur by learning the visual forms of a small vocabulary. Subsequent development of "alphabetic spelling" occurs with the realisation that there are useful regularities in the sound-to-spelling mapping system. When a child has grasped this principle, the ability to spell non-words will emerge, along with an advantage in spelling words which conform to the sound-to-spelling "rules" of English. For the reasons discussed above, however, rules alone are not sufficient. A second "lexical" routine is required to enable the correct spelling of irregular words, and to distinguish between non-homographic homophones such as THEIR and THERE. It is assumed that skilled spellers have both strategies available to them. This "dual-route" approach has provided a successful framework for describing spelling development and classifying developmental and acquired spelling disorders (Ellis, 1984; Frith, 1985; Ellis & Young, 1988).

A consequence of the "dual-route" approach is that word classifications have tended to label items as either "regular" or "irregular," according to whether or not they conform to the sound-to-spelling translation rules of English. However, a connectionist model should be able to exploit regularities at various levels in the mapping from sound to spelling without making use of explicit rules. In the case of reading, the connectionist approach has lead to a different classification of words in terms of the spelling-to-sound "friends" and "enemies" that they possess (Brown, 1987; Jared, McRae & Seidenberg, 1990). According to this approach, only words with enemies would be counted as irregular, but regular words can differ according to the number of friends they have. Table 1 illustrates this

---

The research reported in this paper was supported by grants to the second author from the Medical Research Council (U.K.), the Economic and Social Research Council (U.K.) and the Leverhulme Trust.

classification as applied to the case of spelling, and shows that "number of friends" and "number of enemies" reflect different dimensions (here, as elsewhere, we use the "rime" of a word to classify it: there is ample evidence that this is the relevant psychological dimension: Treiman & Chafetz, 1987).

Word	Friends	Enemies
SOAP	none	HOPE, COPE etc
LOCK	DOCK, ROCK etc	none
BULB	none	none

Table 1: Classification of Word Types

Thus friends and enemies effects can be assessed independently: a difference between "bulb" type words and "soap" type words would reflect an independent effect of number of sound-to-spelling enemies, and a difference between "bulb" type words and "lock" type words would reflect an independent effect of number of sound-to-spelling friends. In the case of single word reading it now appears that there are independent friends and enemies effects on reading time (Brown, 1987, Jared, McRae & Seidenberg, 1990). Many experimental investigations of spelling have, however, not distinguished between sound-to-spelling friends and enemies, and empirically it is not clear which is the relevant factor, although several recent experiments on spelling have found influences of factors other than simple regularity on spelling performance (e.g. Barry & Seymour, 1988). In this paper we report an experimental investigation of this question which focusses on the three word types illustrated in Table 1. First, however, we describe a connectionist model that we have implemented to investigate learning of the English sound-to-spelling mapping system.

A Connectionist Model of Spelling

The model is a feedforward backpropagation net with 50 input, 30 hidden and 50 output units. The training set consisted of 250 words. These included all 63 words used in the experiment discussed below. The remaining words were friends and enemies of words in the group of 63. Input patterns encoded the phonological forms of words, and output patterns represented the corresponding orthographic forms. These patterns were created using a scheme similar to that employed by Seidenberg & McClelland (1989a): words were treated as sets of "triples," (e.g. **SOAP** = **_SO** + **SOA** + **OAP** + **AP_** ; */swp/* = *_sw* + *swp* + *wp_*). Patterns were created independently for all the triples that occurred in the training set. Each triple pattern contained four ON bits and 46 OFF bits: these were chosen quasi-randomly in such a way that similar triples had similar patterns. The pattern for a word was made by superposing the patterns for its triples. The learning rate was set at 0.01, and the momentum parameter was 0.9. The model was given 2000 epochs of learning with the training set.

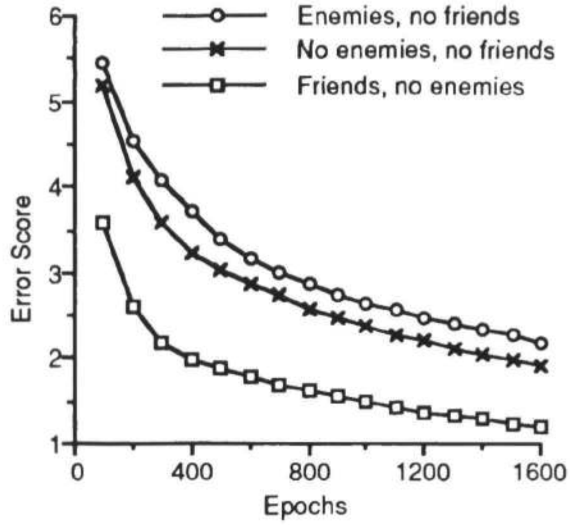

Figure 1: Error Score for 30-Hidden-Unit Model

The vocabulary of the model contained 21 matched word triplets. The words were selected with a view to their later use in the experimental investigation, and corresponded to the three word types in Table 1. Words within each triplet varied in their sound-to-spelling correspondences but were matched as closely as possible on word frequency, positional bigram frequency, and word length. No word in the sample was homophonic with any other English word.

In addition to examining the development of spelling in normal children, we wanted to assess the possibility that the spelling problems experienced by developmental dyslexics could be characterised in terms of reduced computational resources being devoted to the learning process. Similar claims have been made in the case of reading: Seidenberg & McClelland (1989b) found that a version of their spelling-to-sound translation model, which was given fewer hidden units over which to represent its regularities, produced a pattern of performance that was

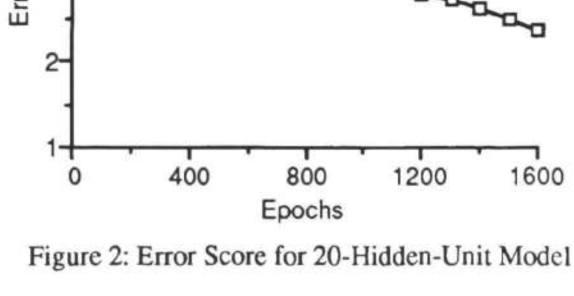

Figure 2: Error Score for 20-Hidden-Unit Model

analogous to that of developmentally dyslexic or delayed readers. In our simulations we adopted a similar approach. The "normal" model was given 30 hidden units, while the "developmentally dyslexic" model was provided with only 20 hidden units during the learning process.

Results of Simulations

The behaviour of the model during learning is presented in Figures 1 through 4. Figure 1 shows the mean summed squared error for each of the three word types during learning in the "normal" model, and Figure 2 shows the same measure for the "developmentally dyslexic" model. In both cases the same ordering of word type difficulty is seen throughout, but the dyslexic model learns more slowly and has higher error scores at any given point in learning.

Of more direct relevance is the actual number of items that the network was able to spell correctly at any point. In order to assess whether a word was spelled correctly, the following procedure was carried out. First, 30 "competitor" spellings were generated for each of the 63 words by randomly changing one of the letters in the correct, "target" spelling. For the words with enemies, one competitor was always the regularised spelling of the word. Thus, the competitors for SOAP included SOPE, SSAP, SOOP, SOAB, TOAP and so on. At each stage of learning, a summed squared error was calculated for all competitors of each word, as well as for the target spelling. A word was classified as correctly spelled, and counted as such in Figures 3 and 4, when the error score for the correct spelling was lower than the error score for *all* of the competitor spellings. Figure 3 shows the *number of words correct* (out of the 21 words of each type) as a function of training for the normal (30 hidden unit) model, while Figure 4 shows the same for the restricted version.

Inspection of Figure 3 reveals that by 2000 epochs of learning the normal network was producing the correct spelling of every word in its vocabulary. At all earlier stages in learning, the model produced the correct spelling of more of the words with friends and no enemies than words with neither friends nor enemies, and even fewer of the words with enemies and no friends.

Comparison of Figures 1 and 3 with Figures 2 and 4 shows that the dyslexic model exhibits a qualitatively similar pattern of performance to the non-dyslexic model, but that performance lags behind the non-dyslexic model at all stages. By the end of 2000 epochs of learning, the dyslexic model correctly spells 20 out of the 21 words with friends only, 17/21 of the words with neither friends nor enemies, and also 17/21 of the words with enemies and no friends. This pattern of selective difficulty with exceptional items is similar to the pattern observed in developmentally dyslexic children.

It is clear that the non-dyslexic model was successful in learning the relevant sound-to-spelling mapping characteristics within its limited vocabulary. Furthermore, the performance of the model leads to a clear prediction about the ordering of word type difficulty during learning.

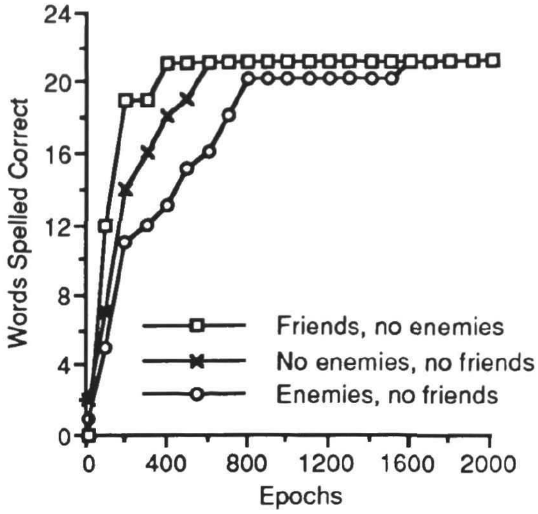

Figure 3: Number of Correct Spellings ("Normal" Model)

The results of the two simulations suggest that the difference between normal and dyslexic spelling development can be well characterised in terms of the amount of computational resources devoted to the task. When insufficient resources are allocated to learning the relevant sound-to-spelling associations, the result is that a lower overall level of performance is achieved at any given stage in learning, but the ordering of the different word types in terms of accuracy is the same. This is consistent with a "delay" rather than a "deviance" picture of dyslexic spelling: the model predicts that when dyslexic and non-dyslexic subjects are matched for overall level of performance, they should show the same qualitative pattern of errors on the three word types. By contrast, "deviance" models assume that qualitatively different processing is involved in developmental dyslexia.

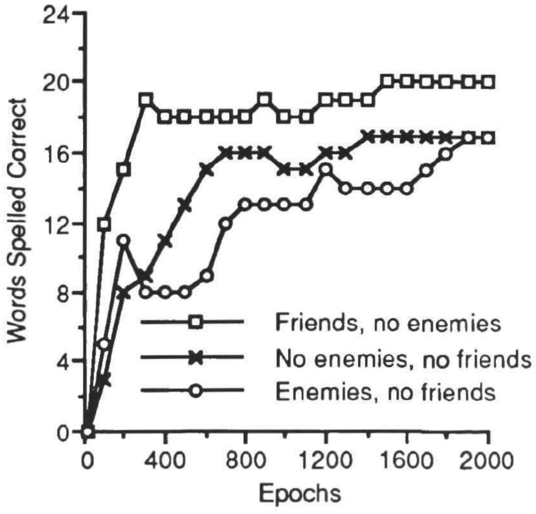

Figure 4: Number of Correct Spellings ("Dyslexic" Model)

Experimental Studies

We now report the results of experiments carried out on both normal and dyslexic children to test the predictions of the models reported above. We describe the spelling performance of four different age-groups of normal children, and a separate comparison between dyslexic and spelling-level matched non-dyslexic subjects.

Subjects

132 non-dyslexic children aged between 91 and 143 months were tested. Data were also collected from 12 "junior" dyslexics and 13 "senior" dyslexics from a specialist school. One of the senior subjects (with the lowest total spelling correct score) was excluded from the analysis to make equal sized groups. All the dyslexic subjects tested in this study had been formally diagnosed as having specific learning difficulties by an independent examiner and were attending a special school for dyslexic boys. Additional tests showed that the junior dyslexics had a reading age 30 months behind their chronological age, and the senior dyslexics were 37 months delayed. Reading ages were assessed by the British Ability Scales.

Materials and Procedure

The stimulus materials for the experiment were the 21 word triplets used in the models. For the spelling test, each stimulus word was presented in a short sentence that used the word in a meaningful context but did not define its meaning. A separate comprehension test was conducted and error rates were examined only for words that were known to individual subjects.

Analysis and Results

The 132 non-dyslexic subjects were divided into four groups of 33 subjects on the basis of spelling ability. To compare normal and dyslexic spellers, control subjects were selected from the sample of normal spellers. They were matched with individual dyslexics on the basis of total correct spelling scores. The dyslexics were on average around two years older than the control subjects spelling at the same level.

The results were examined in terms of the proportion of errors made to *known* words only. Two triplets were withdrawn before analysis because insufficient subjects knew the meanings of one of the words in the triplets. Spelling scores were expressed as the number of correctly spelt words as a proportion of the items known, and subjected to angular transformation prior to analysis. Figure 5 presents the (untransformed) error proportions for the four different age groups of control subjects, and Figure 6 gives the proportions for the junior and senior dyslexics and controls.

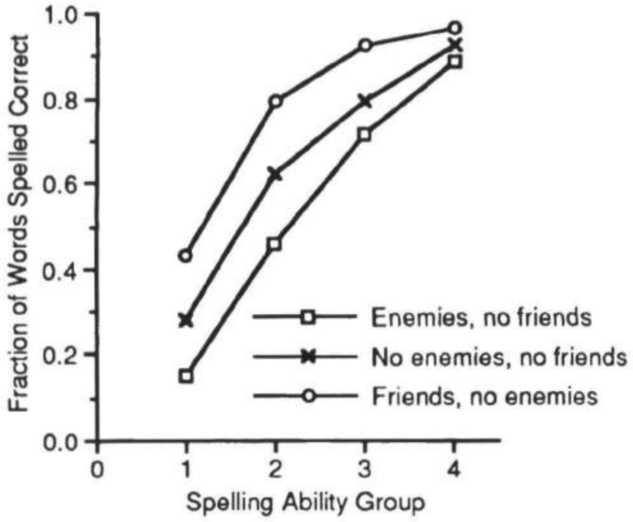

Figure 5: Fraction of Correct Spellings (Normal Subjects)

For normal subjects (data in Figure 5), ANOVA revealed main effects of age-group and word-type, and a significant group X word type interaction reflecting a ceiling effect in the more able groups. Post-hoc analyses of the error rates revealed independent effects of both sound-to-spelling friends and sound-to-spelling enemies within the main effect of word type. ("Friends effects" were assessed by comparing error rates to words with many friends and no enemies to words with neither friends nor enemies; "enemies effects" were assessed by comparing error rates for words with neither friends nor enemies to words with enemies but no friends).

In the error analyses comparing dyslexic and control subjects, there was no significant difference between control and dyslexic subjects (as expected, given that groups were matched on total spelling score) and no significant interactions. There were, however, main effects of both ability group and word type. Post-hoc comparisons of the word type means indicated that there were independent sound-to-spelling friends and enemies effects.

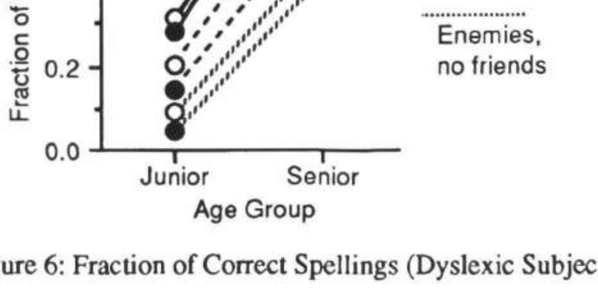

Figure 6: Fraction of Correct Spellings (Dyslexic Subjects)

The results may be summarised as follows. Both dyslexic and control subjects showed independent effects of sound-to-spelling friends and sound-to-spelling enemies on error rates. Furthermore, dyslexic subjects showed the same pattern of effects as younger non-dyslexic subjects matched on overall spelling scores.

These results are in accord with the predictions made by the model. The ordering of difficulty of the three different word types, in terms of the proportion of errors made at different times during development, is exactly the same for the connectionist model and for our human subjects.

Furthermore, despite the fact that we have used a more fine-grained classification of word types than has previously been used to compare normal and dyslexic spelling performance, we find that the dyslexics perform similarly to younger control subjects spelling at the same overall level. This is consistent with the behaviour of our "developmentally dyslexic" model.

Lesioning the Model

We wish to argue that the development of normal spelling, and the delayed development in young dyslexics, can be characterised in terms of the computational demands of mastering a statistical mapping system. A natural question, then, is whether the performance of head-injured patients can be accounted for in terms of a loss of the computational resources available to achieve the requisite mapping (cf. Patterson, Seidenberg & McClelland, 1988).

The relevant phenomena may be summarised as follows. Some patients, generally known as "phonological dysgraphics," lose the ability to spell non-words while the ability to spell real words (whether regular or irregular) is relatively well preserved (e.g. Shallice, 1981). In terms of dual-route models of spelling, this is taken as evidence for loss of the non-lexical sound-to-spelling translation pathway. So-called "deep dysgraphics" exhibit similar problems but also produce semantically-related errors. The complementary syndrome, variously known as "surface dysgraphia" (Ellis, 1984), "lexical dysgraphia" (Beauvois & Derouesné, 1981) or "phonological spelling" (Hatfield & Patterson, 1983) appears to involve a relative sparing of the sound-to-spelling translation routine along with impairment of the lexical spelling routine. These patients therefore have particular difficulty in spelling words with exceptional sound-spelling correspondences. The picture is of course more complex than the simple one presented above (see Ellis & Young, 1988, for a review), and patients vary in the extent of dissociation which they exhibit. For present purposes, however, we are concerned to examine the possibility that our connectionist model of spelling could become "surface dysgraphic" and selectively lose the ability to spell irregular words correctly.

We simulated acquired spelling disorders in the model by setting 2%, 5%, and 10% of the weights in either the input or the output layer to zero. The performance of the model was then examined to see whether it produced the correct spellings of the words in its vocabulary. A total of 10 simulations were carried out at each level of damage.

Results of Lesioning the Model

The results of zeroing 2%, 5% or 10% of the connections in the input or output layers are shown in Figure 7. Performance on words with enemies (i.e. words with exceptional sound-spelling patterns) does indeed drop off more quickly with increasing damage. Thus with 10% damage to the connections between input and hidden layers, the model still correctly spells 13 out of the 21 friends-only words, but only 8 of the 21 words with enemies.

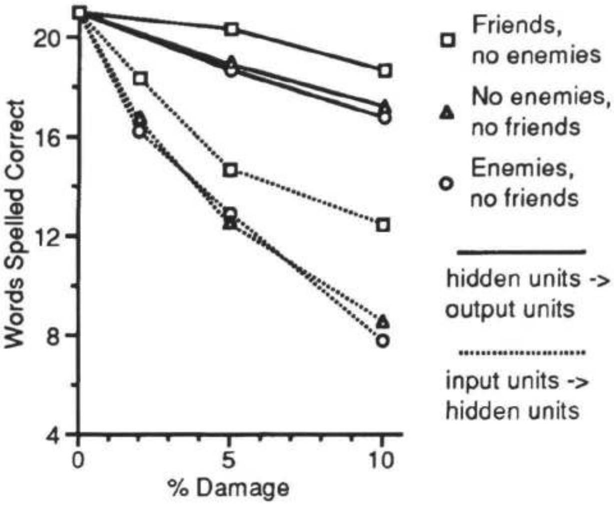

Figure 7: Performance of the Model After Lesioning

How does this compare with the pattern of performance produced by surface dysgraphic patients? Some patients show a sharper dissociation than our model: Shallice (1988) summarises the results of several studies on the relevant patients. Most patients show very good performance on non-words and words with low ambiguity (i.e. regular items), performing at over 90% on these categories. Performance on words with medium levels of ambiguity ranged from about 70% to 100%, and performance on highly ambiguous items varied from 40% to 90%. Thus only a very few patients show a sharp dissociation, although all perform worse on the ambiguous items.

The average performance of our model at low levels of damage was similar to the relatively unimpaired patients, but Figure 7 shows that a 40% level of performance on our exception words (those with enemies) was accompanied by just over 60% correct performance on words with friends only. This appears to contrast with the sharper dissociations exhibited by some of the patients, but two further factors need to be taken into consideration. First, Figure 7 represents the average of 10 lesioning simulations. There was considerable variability in the simulations, as in the patient populations, and some of our individual runs showed more extreme dissociations. In several simulations no friends-only words were lost but 5 exception words were lost. In one run, after just 2% damage, 20/21 of the regular words (with friends) were spelled correctly but only 14/21 of the exceptional items. A further point is that the vocabulary of our model was relatively small, and so it included only a few enemies for our exception words. The patient data shows that much less severe impairments are found with words of only medium level ambiguity, and this perhaps provides the best characterisation of the words in our model, given its limited vocabulary. With this assumption, the model's performance provides a good fit to the empirical data. Furthermore, the inclusion in our simulation of words with neither friends nor enemies leads to a prediction (untested to our knowledge): surface dysgraphics will have difficulty with words with neither friends nor enemies as well as with "exception" words, even though the words with neither friends nor enemies are not ambiguous in their sound-to-spelling correspondence.

Conclusions

We have attempted to use the connectionist approach to provide a computationally explicit characterisation of one aspect of human spelling development. We should note that we see the model as complementing rather than replacing the descriptive-level insights from developmental and cognitive psychology.

The model is successful in that its performance closely mirrors that of human subjects under a wide range of circumstances. We conclude that much of the process of learning to spell is usefully viewed as a statistical one. Observed performance in both developmental and acquired disorders can be characterised as being due to a lack of, or loss of, computational resources over which to represent the statistical regularities and sub-regularities implicit in the English sound-to-spelling mapping system. This abstract, high-level characterisation is, of course, far removed from any strictly neurobiological interpretation of the connectionist approach. However some researchers, (e.g. Galaburda, 1989), have pointed to the parallels between connectionist characterisations of developmental reading disorders in terms of provision of computational resources, and the atypical symmetry distribution of the *planum temporale* found in dyslexic brains.

Our model is clearly limited in a number of respects. It operates with a small and constant vocabulary, and contains no mechanism for modelling early "logographic" spelling or strategic change. Rather, it is concerned only with the development of *alphabetic* spelling, and so does not address the issue of developmental stages and strategies. Furthermore, it only accounts for one type of acquired spelling disorder: extensions to the model would be necessary to account for complementary syndromes such as phonological dysgraphia, and also to enable the model to distinguish non-homographic homophones such as THEIR and THERE. A final comment concerns the nature of the "triple" representation used on both input and output units. As Prince & Pinker (1988) have shown, such representations are not adequate for representing the whole of English vocabulary, although they suffice for the 250 word vocabulary we have modelled to date.

We believe that these limitations do not compromise our main conclusion, however. This conclusion is that much psychological data concerning normal spelling development, developmental dyslexia, and acquired spelling disorders can be well described in terms of the degree to which the various subject populations have mastered the statistical properties of the English sound-spelling mapping system.

References

Barry, C., and Seymour, P. H. K. 1988. Lexical Priming and Sound-to-Spelling Contingency Effects in Nonword Spelling. *Quarterly Journal of Experimental Psychology* 40(A):5-40.

Beauvois, M-F, and Derouesné, J. 1981. Lexical or Orthographic Dysgraphia. *Brain* 104:21-50.

Brown, G. D. A. 1987. Resolving Inconsistency: A Computational Model of Word Naming. *Journal of Memory and Language* 23(1):1-23.

Ellis, A. W. 1984. *Reading, Writing and Dyslexia.* London: LEA.

Ellis, A. W., and YOUNG, A. W. 1988. *Human Cognitive Neuropsychology.* London: LEA.

Frith, U. 1985. Beneath the Surface of Developmental Dyslexia. In K. E. Patterson, J. C. Marshall and M. Coltheart, eds. *Surface Dyslexia.* Hillsdale, NJ: LEA.

Galaburda, A. M. ed. 1989. *From Reading to Neurons.* Cambridge, Mass.: MIT Press.

Hatfield, F. M., and Patterson, K. E. 1983. Phonological Spelling. *Quarterly Journal of Experimental Psychology* 35(A):451-468.

Jared, D., McRae, K., and Seidenberg, M. S. 1990. The Basis of Consistency Effects in Word Naming. *Journal of Memory and Language* 29:687-715.

Patterson, K. E., Seidenberg, M. S., and McClelland, J. L. 1988. Connections and Disconnections: Acquired Dyslexia in a Computational Model of Reading. In R. G. M. Morris, ed. *Parallel Distributed Processing: Implications for Psychology and Neurobiology.* Oxford: OUP.

Prince, A., and Pinker, S. 1988. Wickelphone Ambiguity. *Cognition* 30:189-190.

Seidenberg, M. S., and McClelland, J. L. 1989a. A Distributed, Developmental Model of Word Recognition and Naming. *Psychological Review* 96:523-568.

Seidenberg, M. S., and McClelland, J. L. 1989b. Visual Word Recognition and Pronunciation: A Computational Model of Acquired Skilled Performance and Dyslexia. In A. Galaburda, ed. *From Reading to Neurons.* Cambridge, Mass.: MIT Press.

Shallice, T. 1981. Phonological Agraphia and the Lexical Route in Writing. *Brain* 104:413-429.

Shallice, T. 1988. *From Neuropsychology to Mental Structure.* Cambridge: CUP.

Treiman, R., and Chafetz, J. 1987. Are There Onset- and Rime-Like Units in Written Words? In M. Coltheart ed. *Attention and Performance XII.* London: LEA.