Brown, G.D.A. & Loosemore, R.P.W. (1995). A computational approach to dyslexic reading and spelling. In R. M. Joshi & C. K. Leong (Eds.), *Developmental and acquired dyslexia: Neuropsychological and neurolinguistic perspectives.* Dordrecht: Kluwer Press.

A computational approach to dyslexic reading and spelling

Gordon D. A. Brown
Richard P.W. Loosemore

Cognitive Neurocomputation Unit,
Department of Psychology,
University of Wales, Bangor
Gwynedd LL57 2DG
United Kingdom
Earn/BITNET: PSSØØ1@uk.ac.bangor

Abstract

This chapter discusses the ability of computational models to improve our understanding of dyslexic reading and writing. Connectionist models of the development of alphabetic reading and spelling in normal and dyslexic children are described. The models learn to associate representations of word pronunciations with spellings. The models learn to read and spell regular words more quickly than irregular items. When the computational resources available to such models are restricted, the models learn more slowly and fail to learn some of the irregular items in their vocabularies. The restricted models behave analogously to developmental dyslexics, and, crucially, have selective deficits with non-word processing although they do not show reduced sound-to-spelling or spelling-to-sound regularity effects. This is consistent with the experimental literature. Experimental evidence is reported that shows that both normal and dyslexic children of various ages have difficulties with reading and spelling particular word types that are similar to the problems experienced by the models on the same words. 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. The level of difficulty posed by particular words in spelling is well predicted by the extent to which those words conform to the relevant regularities of the language. Furthermore, the models resolves an apparent paradox in the experimental literature, for in their dyslexic forms they exhibit a selective deficit in non-word reading and spelling even though they do not show reduced sound-to-spelling regularity effects.

Keywords

Children; Computational Models; Connectionism; Development; Reading; Regularity; Spelling

Introduction

In this chapter we show how recent advances in computational modelling can improve our understanding of the cognitive deficits associated with developmental dyslexia. In particular we focus on the nature of the spelling process in normal and dyslexic children, and describe some of our recent experiments that have studied spelling error rates in dyslexia. We argue that computational modelling can resolve apparently contradictory research findings in the dyslexia literature on both reading and spelling. The plan of the chapter is as follows. First, we discuss the demands of the spelling process and current psychological models of the normal development of spelling. We then describe the attempts that have been made to characterise the nature of the spelling deficit in dyslexia in terms of these cognitive psychological models. In the next part of the chapter we go on to describe the class of computational models called "connectionist" or "parallel distributed processing" models. We illustrate this approach with particular reference to a recent connectionist model of the development of reading (Seidenberg & McClelland, 1989a; 1989b) and show how this model provides an alternative way of looking at dyslexics' problems in reading (Seidenberg, 1989). We also describe our own investigations of the performance of a similar connectionist model of reading, and argue that the connectionist approach sheds new light on empirical data that are otherwise difficult to interpret. Next, we describe a connectionist model of spelling that we have developed (Brown, Loosemore & Watson, 1992) and describe the model's predictions concerning the types of word that should be particularly difficult for normal and dyslexic children to spell. We then summarise the results of some experiments that test the predictions of the model as applied to developmental dyslexia. The spelling model is analysed in detail, and it is shown that "dyslexic" versions of it exhibit a selective deficit in non-word processing. This is consistent with recent experimental work, which has generally found selective deficits in non-word processing in dyslexia even though reductions in regularity effects are not normally apparent in group studies. In the final part of the chapter we discuss the implications of the approach for our understanding of normal and dyslexic spelling development. We argue that the connectionist approach is a useful one and that it can provide a novel way to address some theoretical issues in the domain of dyslexia research. In particular we conclude that much of dyslexics' observed performance can be characterised in terms of a shortfall in the computational resources available for reading and spelling acquisition.

We begin with spelling, as this forms the main focus of the present chapter. However many of the conclusions apply to reading as well as spelling, and we shall try to bring this out throughout the chapter.

The nature of spelling

Many English words can be correctly spelled on the basis of their pronunciation. However, the use of pronunciation alone is an unreliable method of deriving spellings. Some words are pronounced the same as other words, but spelled differently (non-homographic homophones: HARE-HAIR; THEIR-THERE). Other words, such as SOAP, are not spelled as might be expected from their pronunciations - cf. HOPE, COPE, ROPE, etc. It is of course theoretically possible that pronunciation information is not used at all in spelling. There is, however, ample empirical support for the claim that pronunciation information is used in both children's and adults' spelling. Any satisfactory model must reconcile these two sets of constraints.

Psychological models of spelling

Descriptive approaches to spelling development often assume that the early ("logographic") stages of learning to write involve acquiring the visual forms of a small number of items. In this stage sub-lexical structure is not used. In subsequent development the child gradually becomes aware of this structure and uses it to develop a sound-to-spelling translation routine. This is "alphabetic" spelling. Such a routine can cope with regular but not irregular words. Finally comes what Frith (1985) terms "orthographic" processing: this involves the "instant analysis of words into orthographic units without phonological conversion" (1985, p. 306).

Skilled readers thus have alternative strategies for spelling words available to them. One routine makes use of sound-to-spelling translation "rules" of some kind, although there is considerable debate about the representations that these rules operate on (e.g. Barry & Seymour 1986; Campbell, 1985). The "direct" routine provides a one-to-one mapping from particular lexical entries to representations of their spelled forms. This routine is used for words which cannot be reliably spelled on the basis of their pronunciations. Recent models typically contain other components such as a graphemic output buffer (e.g. Caramazza, Miceli, Villa & Roman, 1987).

Frith (1985) suggests that classical developmental dyslexia can be characterised as an arrest at her "stage 1" during which writing is logographic. Sound-to-spelling translation routines do not develop, except perhaps as a result of careful individual tuition, and the child is left with a mechanism that can only deal with words as wholes, and which is not sensitive to sub-lexical regularities. This view of developmental dyslexia as arrest at the logographic stage leads to a number of predictions. For example, the lack of translation routines should lead to a selective difficulty in reading or spelling non-words, and also to reduced or absent spelling-to-sound and sound-to-spelling regularity effects. (This is because such regularity effects reflect the use of the sub-lexical translation routines: regular words will only be advantaged if sound-to-spelling knowledge is available). Later in the present chapter we describe a study we designed (Brown, Loosemore & Watson, 1992) to investigate whether dyslexics show a normal regularity effect in spelling. This topic has not been widely studied in spelling: Barron (1980) found larger sound-to-spelling regularity effects in poor readers, and Seymour and Porpodas (1980) report data which may suggests a smaller regularity effect for dyslexic children on low frequency words. In spelling, as in reading, there is already evidence that dyslexics do have particular difficulty with non-word processing (e.g. Frith, 1980; Jorm, 1981). It should be noted that this approach views dyslexic processing after the time of developmental arrest as abnormal, or "deviant," rather than merely "delayed." We address this issue in more detail below.

The traditional psychological approach outlined above makes frequent reference to "cognitive-level" concepts such as rules, strategies and developmental stages. The framework has proved successful in characterising both normal and disturbed literacy development, and has lead to much fruitful research. However our own interest is in implementing computational models of these processes. It is our belief that, as Kelvin put it: "I can't really understand something unless I can make a mechanical model of it." Further: we would like to produce a model which *acquires* knowledge of language, to see whether the learning process itself can result in the observed characteristics of skilled spelling. In trying to translate the above models into working implementations, we find that they lack specificity in precisely those aspects needed to produce models which learn. Their strengths lie in their descriptive coverage of the data, rather than in their provision of a low-level causal account of the mechanisms that mediate the acquisition of literacy.

Connectionism

An alternative approach is provided by "connectionist" or "parallel distributed processing" models of psychological processes. The emphasis here is not on high-level concepts such as rules or strategies; rather, such models can learn to associate pairs of patterns without reference to explicit rules. For example, the recent model of Seidenberg and McClelland (1989a) learns to associate the orthographic forms of words to their corresponding phonological forms.

A connectionist network consists of a large number of computational units whose behaviour is in some respects akin to that of neurons. Each unit has connections to some, although generally not all, of the other units in the network. Associated with every unit is a quantity called the “activation” level of that unit. The main purpose of a connection is to communicate the activation level of one unit to another. The activation level is modified as it passes through a connection, depending on the “strength” of that connection. The level of activation of each unit is determined by the sum of all the modified activations that it receives from other units that it is connected to. Broadly speaking, a unit’s activation level is proportional to the amount that the total activation coming into the unit gets over a “threshold” value. The amount of influence that one unit has on another (if they are connected) depends on both the size of its own activation level, and the strength of the connection between the two units.

Such networks have a number of interesting computational properties. Some units in the network can be considered “input” units, while others are “output” units. The activation levels of the set of input units can be thought of as a pattern which represents something meaningful, such as the phonological form of a word. Likewise, the pattern formed by the activation levels of the output units may be regarded as a representation of, for example, the spelled form of a word. When an input pattern is imposed on the input units, activation will spread through the connections in the network until some pattern of activity is established on the output units. The precise pattern formed on the output units will depend on the strengths of all the connections in the network. The fundamental property of a connectionist network is that given a particular input pattern, the connection strengths can change themselves in such a way as to cause a particular pattern to appear on the output units in response to an input pattern. In this way, the network can learn to associate one pattern with another. Moreover, a network is not confined to representing just one input-output pair; rather, a number of associations between pairs of patterns can be learned within the same set of connections.

It is important to emphasise that the network learns about associations between pairs of patterns simply by being repeatedly exposed to the pairs: all of the learning takes place by slow modification of connection strengths in the course of “experience.” At the end of learning the network is able to produce the correct output pattern in response to a particular input.

It is this learning ability which has led to a widespread interest in these models within psychology, for it is possible to examine the performance of the network as it encodes a set of associations over time. Connectionist networks have been used to provide psychologically interesting models of a variety of different behaviours, such as verb tense learning, speech perception and speech production. In the present chapter we focus on the application of connectionist techniques to the modelling of literacy development.

Early network models of reading (Brown, 1987; McClelland & Rumelhart, 1980) lacked any ability to learn associations for themselves. This limited the size of vocabulary they could work with, for all the connection strengths had to be determined "by hand." One more recent and highly influential model of reading has been developed by Seidenberg and McClelland (1989a). This model works using the principles outlined above - it learns to associate input representations of word orthography with output representations of word pronunciations. Using a standard connectionist learning procedure it learns to produce the correct pronunciation of nearly 3,000 monosyllabic words. Seidenberg and McClelland (1989a) show how this model can account for a very wide range of psychological data from a variety of different experimental paradigms such as lexical decision tasks and single word naming. The model learns to abstract some of the general statistical regularity and redundancy that is present in the relationship between orthography and phonology in English. Indeed, this abstraction of statistical structure is an important general characteristic of connectionist models. This seems to underlie the ability of the model to show spelling-to-sound regularity effects which are very similar to those observed in human performance. These regularity effects are more pronounced for low frequency words, which is the case for human subjects (Seidenberg, Waters, Barnes & Tanenhaus, 1984). In addition, the model is sometimes able to synthesise appropriate pronunciations for novel items (non-word naming). Furthermore, Patterson, Seidenberg and McClelland (1989) have argued that when the model is "lesioned" by removing some proportion of its units or connections, its behaviour resembles in many respects that of brain-injured "acquired dyslexic" patients, in that it has a selective difficulty in reading words which contain irregular spelling-to-sound correspondences. Here we do not summarise the full range of phenomena encompassed by this model, nor do we discuss some recent criticisms of the model (e.g. Besner, Twilley, McCann & Seergobin, 1990).

We do, however, wish to focus on two particular properties of the model. The first is its ability to pronounce both regular and irregular words with only one mechanism - it therefore stands in contrast to so-called "dual-route" models of reading which, analogously to the psychological models of spelling described earlier, assume that two separate routes to pronunciation must be available if both regular and irregular items are to be pronounced successfully. (Note that the full architecture described by Seidenberg & McClelland, not all of which was implemented, does contain two routes.) The second claim that has been made on the basis of this connectionist model of reading concerns its ability to characterise the reading behaviour of developmentally dyslexic children in terms of the computational capacity made available to the network during learning.

To understand how this works, it is necessary to understand that not all of the units in a connectionist network need be either “input” or “output” units. Those which are neither input nor output are called “hidden” units. The model appears to use its hidden units to represent regularities in the corpus of patterns that it sees. Because it is not generally provided with sufficient hidden units to enable it to learn all the required associations on a one-to-one basis, it must choose economical representations such that it can encode many patterns over a few units. The capacity of the model to do this will depend upon the number of hidden units available. The connectionist model that Seidenberg and McClelland used to examine skilled adult reading was given 200 hidden units. Providing the model with only 100 hidden units (reducing its computational capacity) resulted in a general reduction in performance for all word types (regular and irregular; high and low frequency). The model with reduced computational resources showed spelling-to-sound regularity effects for both low frequency and high frequency words, whereas the larger model, like skilled adult readers, only showed a regularity effect for the low frequency words. Seidenberg and McClelland (1989b) argue that this is similar to the difference between good and poor young readers, for poorer readers show spelling-to-sound regularity effects for both high and low frequency words, whereas good readers only show regularity effects for low frequency items.

It is therefore possible to argue that poor reading can at least partly be described in terms of a limitation in the computational resources available to the model during learning. Note that this represents a different way of looking at reading problems from the traditional debate, which is couched in terms of whether reading is “deviant” or “delayed.” The behaviour of the model with restricted computational resources cannot adequately be described as simply “delayed” or “deviant.” One way of characterising the delay hypothesis of reading disorders is to say that dyslexics read “in the same way as” younger non-dyslexics. In this respect the model appears to conform to a delay view of reading disability. However, it is also the case that the dyslexic model shows regularity effects when the non-dyslexic model does not (i.e. on high-frequency words). In the traditional information processing framework the presence of regularity effects in one group but not another would be taken as evidence for qualitatively different (i.e. “deviant”) processing. More specifically, the presence of regularity effects would be taken to indicate a failure to move from alphabetic to orthographic processing. Yet it appears nonsensical to interpret the dyslexic model as employing a qualitatively different processing strategy simply because it has fewer computational resources available to it. It therefore seems that the mere presence or absence of spelling-to-sound regularity effects cannot be taken as evidence for or against the use of a particular processing strategy. We return to discussion of this issue below.

It is not of course claimed that the connectionist approach and the variations in network capacity can account for the whole range of dyslexic symptoms. Nevertheless, in terms of the regularity effects, which provide one of the main indicators of which stage a child is at in reading, the connectionist model of reading can account for a wide range of the relevant empirical data.

Reading, connectionism and dyslexia

An obvious further question is, therefore, whether the connectionist approach to reading development can account for the full range of experimental evidence from the study of developmental dyslexia. In the case of reading as for spelling, the majority of experimental work has been directed towards testing the “delay” and “deviance” accounts of dyslexia. As we described earlier in the context of spelling, the phonological deficit hypothesis leads to two critical predictions about experimental tasks that should cause particular difficulty for dyslexic children in reading. The first of these concerns the reading of non-words. Because a non-word (e.g. SLINT) can only be pronounced by using spelling-to-sound rules or analogies, non-word reading provides a test of spelling-to-sound decoding ability. Therefore, if dyslexics have specific decoding problems, they should perform worse at non-word reading than non-dyslexics. This prediction can be tested by comparing the performance of dyslexic children with younger non-dyslexic children who are reading at the same level. (The rational behind this design is that it allows tests of whether a given level of reading skill is achieved by the same or different strategies in different populations. If dyslexic and non-dyslexic children of the same *chronological* age were compared, in contrast, then any group differences could be due to the smaller amount of reading practice experienced by the dyslexic children - differences could be a *consequence* rather than a *cause* of the reading delay.)

Rack, Snowling and Olson (1992) review considerable evidence that developmental dyslexics do have difficulty in non-word reading or repetition when compared with control subjects reading at the same level (e.g. Bradley & Bryant, 1981; Frith & Snowling, 1983; Seymour & Porpodas, 1980; Snowling 1981; Snowling, Goulandris, Bowlby & Howell, 1986). However we have argued elsewhere that the majority of studies that have looked for reduced effects of spelling-to-sound regularity in reading have found equivalent regularity effects in dyslexics and controls (e.g. Brown & Watson, 1991). Thus several studies have found equivalent spelling-to-sound regularity effects in dyslexic reading (e.g. Baddeley, Logie & Ellis, 1988; Beech & Harding, 1984; Brown & Watson, 1991; Seidenberg, Bruck, Fornarolo & Backman, 1985; Szesulski & Manis, 1987; Treiman & Hirsh-Pasek, 1987; Watson & Brown, 1992) but relatively few have found reduced regularity effects (Barron, 1980; Frith & Snowling, 1983). This represents a somewhat contradictory set of findings, for if dyslexics are indeed impaired at alphabetic processing one should find both non-word difficulty *and* reduced regularity effects in both reading and spelling.

To explore this paradox further, we examined regularity effects and non-word processing in a connectionist model of reading (Brown, Watson & Loosemore, 1993). This was essentially a smaller and simplified version of the model described by Seidenberg and McClelland (1989a). The network learned to associate word pronunciations with corresponding orthographic representations using the "backpropagation" gradient descent learning algorithm (Rumelhart, Hinton & Williams, 1986), as in the Seidenberg and McClelland model. Words were represented as activations of "triples" of phonemes or letters. The orthographic form of a word such as HAVE, for example, would be represented by the four triples _HA + HAV + AVE + VE_ (where the "_" character signifies a word boundary). Although this scheme cannot represent all words (Prince & Pinker, 1988) it suffices for the vocabulary of our model. Choice of representational scheme may not be critical provided sufficient sub-lexical structure is captured. Pronunciations were represented as distributed patterns of activation over 50 "output" units - each output unit participated in the encoding of 24 of the phoneme triples. Orthographic patterns were represented over 50 "input" units in a similar way. The input units and output units were connected via an intermediate layer of "hidden units." The number of hidden units was varied in order to vary the computational resources of the network in an attempt to model dyslexic performance. There were three versions of the model: "normal" (35 hidden units); "mildly dyslexic" (20 hidden units), and "severely dyslexic" (15 hidden units). The model was given a vocabulary of 19 regular words, 19 irregular words, and 189 other words selected so as to render the critical items regular or irregular for the model. For example, the irregular word PINT was accompanied in the vocabulary by enemies MINT, HINT, and TINT. The performance of the model was assessed in terms of the "summed squared error score" to each item as it changed during learning. This error measure represents the difference between the target and the actual pronunciation for each association the network was required to learn.

The results suggested that the early stages of learning in a model with relatively high computational resources are qualitatively similar to later learning in the models with restricted resources. Thus when the behaviour of the network is assessed in terms of its performance on irregular and irregular items, the "dyslexic" versions of the model show delayed rather than deviant performance - the models all showed the same relative difficulty on the different item types (irregular items giving rise to a higher error score), but the "dyslexic" versions of the model, with reduced computational resources, learned more slowly. This is consistent with the results of the experimental evidence described above, which has generally found equivalent-sized regularity effects in dyslexic children and younger non-dyslexic children.

However, we noted above that empirical research in dyslexia has reached a different conclusion when non-word reading is used as the performance measure. We therefore examined non-word processing in the three versions of the reading model. Non-words were derived from each regular and irregular word by changing one of the word's consonants (e.g. YILL; MAVE). Non-word performance was assessed by presenting the model with phonological representations of non-words and examining the error score for the "correct" (regular) pronunciation of the non-words. Differences in regularity effects and non-word processing were then examined in the dyslexic and non-dyslexic models. An analogue to a reading-age control experiment was carried out by taking the three models at the point in learning at which they all had an equal error score to regular words. This point was reached after 120 epochs of learning for the normal model, and after 345 and 1200 epochs of learning for the "mildly dyslexic" and "severely dyslexic" models respectively. The different models can therefore be considered to be matched on "reading age" rather than "chronological age" for the purposes of this comparison.

INSERT FIGURE ONE ABOUT HERE

Figure One illustrates the critical result. It shows that the two "dyslexic" models showed equal spelling-to-sound regularity effects to the "non-dyslexic" model but greatly increased error scores to non-words (whether these non-words were derived from regular or irregular words). This is the pattern of results observed in experimental studies of dyslexia, which have shown that dyslexics have more difficulty than reading-age-matched non-dyslexic children in reading non-words, even though they show equivalent regularity effects. The model therefore behaves in a way which is paradoxical when interpreted in terms of the "phonological deficit" hypothesis described earlier, for the non-word processing deficit has previously been taken to reflect "deviant" processing in dyslexia, while the equivalent-sized regularity effects have been taken to reflect "normal but delayed" processing in dyslexia.

The results must be treated with some caution owing to the restricted size of the model's vocabulary. Nevertheless, they demonstrate that the relative performance of the models with differing computational resources depends upon the performance metric that is adopted. Examining the difference between the models' error to regular and irregular members of the set to be learned does not show differences between systems with differing computational resources, while examining error to novel items does. It appears that non-word processing is a more sensitive measure of the generalisation capacity of a reading system than is the regularity effect. When computational resources are restricted in the model, it will use the resources that are available to it to learn the words in its vocabulary, and any residual computational capacity will be used for generalisation. We return to these issues in the general discussion.

A connectionist approach to spelling

We now describe our own connectionist model of spelling development (Brown, Loosemore & Watson, 1992), in an attempt to see if the connectionist approach can also account for some of the observed phenomena in this area. Spelling is a more difficult process than reading, computationally speaking, because the mapping from phonology to orthography in English is more irregular and ambiguous than the reverse process. Nevertheless, connectionist models have already exhibited some success in learning to spell (Olson & Caramazza, 1988).

How the model works

The model is a three-layer feedforward network similar to the reading model described above. The input and output layers both contain 50 units, and the hidden layer of the “non-dyslexic” model version has 30 units. Fuller details can be found in Brown, Loosemore & Watson (1992). Each input pattern represents the phonological form of a word, while the corresponding target output pattern represents the orthographic form. In the simulations reported here the network was trained to spell a set of 223 single-syllable words. The model learns the associations between pronunciations and spellings using backpropagation, as in the models of reading described above. This simply involves repeatedly adjusting the strengths of all the connections in the network, a little at a time, in such a way that the performance of the network gradually improves over time (Rumelhart, Hinton & Williams, 1986).

Vocabulary

Our main interest in the current modelling enterprise is with the model’s ability to spell regular and irregular words. Irregular words have been defined as those which do not conform to the sound-to-spelling “rules” of English. However, an alternative explanation for the apparent effects of *regularity* may be given in terms of sound-spelling *friends* and *enemies*. In the remainder of this paper we will continue to use the terms “irregular” and “regular,” as this is most consistent with current usage, but “irregular” words will be taken to be those with only enemies, while “regular” words will be those with only friends.

It was necessary to devise a vocabulary that could be used both as input to the connectionist model and for use in the experiments on normal and dyslexic children described below. Nineteen pairs of words were produced and each pair contained a regular and matched irregular item. (We also examined a third class of word, but these results are not discussed here.) The words in each group 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. The experimental set consisted of 19 such pairs. The network model learned a total of 223 words. The remaining words were included in order to give the regular and irregular words some friends and enemies respectively. For each regular word there were, on average, four words with similar orthography and phonology to act as friends, while each irregular word had an average of four words with similar phonology, but different orthography, to act as enemies.

A distributed encoding scheme similar to that employed by Rumelhart and McClelland (1986) and Seidenberg and McClelland (1989a) was used to create both the input and target patterns, in such a way that there was a tendency for phonologically similar words to have similar input patterns, and for orthographically similar words to have similar target patterns.

Assessment of performance

The pattern error score (sum of the squares of the errors at the output units) of the actual output pattern with respect to the target output pattern was used as a relative measure of spelling accuracy, for comparing performance on different types of words. This measure is straightforward of interpretation: the higher the error score, the greater the difficulty the model has with learning that spelling and the less likely it would be that a correct spelling could be produced.

Developmental dyslexia in the model

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, as we argued above in the case of reading. In our simulations of spelling we adopted a similar approach. The “normal” model was given 30 hidden units, while a “mildly dyslexic” model was provided with only 20 hidden units during the learning process, and a “severely dyslexic” model was given only 15 hidden units.

Results of simulations

Word spelling in the model

For all models, irregular words had the highest error score, and regular words had the lowest error score. Figure 2 shows the error score for regular and irregular words in the three different versions of the model. All show sound-to-spelling regularity effects, revealed in higher error scores for the irregular items, but the "dyslexic" versions of the model learn more slowly and never achieve the same level of accuracy as the non-dyslexic versions.

INSERT FIGURE 2 ABOUT HERE

In order to assess non-word performance in the model we derived "regular" and "irregular" non-words based on the regular and irregular words. Each non-word was created by replacing the onset phoneme cluster in the phonological form of the word (e.g. */swp/* -> */fwp/*). An input pattern based on this non-word could then be presented to the input layer, and the resulting pattern at the output compared with a pattern representing the target spelling of the non-word. In the case of non-words derived from irregular words, the target spelling was the regular form.

The three models, non-dyslexic (35 hidden units), mildly dyslexic (20 hidden units) and severely dyslexic (15 hidden units) were matched on their performance in spelling regular words. We examined error score for non-words and irregular words when the three different models showed an equal error score on the regular items. Thus this is a spelling-level match - the non-dyslexic model reached this level of performance after 130 epochs of learning, the mildly dyslexic model after 390 epochs, and the severely dyslexic model did not reach this level of performance until 1580 epochs of learning.

INSERT FIGURE 3 ABOUT HERE

The results can be seen in Figure 3. It can be seen that the dyslexic models show almost equal error scores for the irregular items, but they show a dramatic rise in error scores for non-words derived from consistent words, and non-words derived from irregular words, as the number of hidden units is reduced.

How should we interpret these results? Non-word processing ability and regularity effects have been seen as two different indicators of the presence of alphabetic processing. These measures have lead to contradictory theoretical interpretations in the literature. We suggest that this is because non-word processing ability provides a more sensitive measure of the generalisation capacity of a system. Thus experiments which look for non-word processing deficits in dyslexics compared with ability-matched controls are far more likely to find deficits than are studies which look for reduced sound-to-spelling regularity effects in dyslexic populations.

The results of the model which simulates dyslexia by using fewer hidden units are taken to 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.

Experimental Studies

In this section we describe the results of an experiment that we carried out on normal adults and dyslexic children to test the predictions of the model.

Sound-to-spelling regularity in dyslexia

The connectionist model of spelling that we described briefly above predicts that dyslexic children should show equivalent sound-to-spelling regularity effects when compared with non-dyslexic children reading at the same level. We tested this prediction by examining the spelling performance of 24 dyslexic and 24 matched non-dyslexic children. 12 dyslexic children came from a Junior class, and 12 from a Senior class. All the dyslexic subjects had been formally diagnosed as having specific learning difficulties by an independent examiner and were attending a special school for dyslexic boys. Additional tests (using the British Ability Scales) showed that the Junior dyslexics had a reading age 30 months behind their chronological age, and the Senior dyslexics were 37 months delayed.

The stimulus materials used in the computational models were also those used in the experiment. 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. We also conducted a separate comprehension test, and analysed error rates only to words that were known to individual subjects.

Figure 4 presents the (untransformed) error proportions for the Junior and Senior dyslexics and controls.

INSERT FIGURE 4 ABOUT HERE

In the comparison between dyslexic and control subjects there were main effects of both ability group and word type. There was no significant difference between control and dyslexic subjects (as expected, given that groups were matched on total spelling score) and, crucially, there were no significant interactions. Thus, we found 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, which is provided with fewer hidden units over which to represent the statistical regularities inherent in the sound-to-spelling mapping problem.

The fact that our dyslexic and non-dyslexic children showed regularity effects of equal magnitude fails to support the hypothesis that the dyslexic children have not attained a stage of alphabetic processing, at least in spelling. Of course, these results do not exclude the possibility that the dyslexic subjects may differ in *reading* strategies: it is entirely possible that the dyslexics attain alphabetic spelling but not alphabetic reading, although the numerous studies (cited earlier) which have failed to find reduced spelling-to-sound regularity effects in dyslexic reading go against this conclusion. We should also note that our experiments militate against the conclusion that dyslexics show an *over*-reliance in phonological coding during spelling (Barron, 1980), for if this were so the dyslexics should show larger regularity effects, and in our experiment they did not do so.

We now examine the implications of the results for a number of the theoretical issues that were raised at the beginning of this chapter.

Discussion

Delay or deviance?

We now assess the extent to which connectionist models can illuminate some critical theoretical issues in the study of developmental dyslexia. As we described in the introduction, one important question has been whether the processing of dyslexic children is “delayed” or “deviant.” The most common version of the “deviance” hypothesis is that dyslexics do not progress to an “alphabetic” stage of reading and spelling in which they make fluent use of sound-to-spelling or spelling-to-sound rules. This does not preclude the possibility that they may, after the point of developmental arrest, go on to develop compensatory strategies of some kind (perhaps as a result of instruction, or perhaps involving greater development of a visual/lexical non-alphabetic spelling routine). The "delay" hypothesis, in contrast, implies that dyslexic children progress through the same stages as non-dyslexic children, but at a slower rate.

In the study of both reading and spelling, two different experimental strategies have been used to determine whether or not dyslexic children who are processing at the same overall level as control groups are processing in a qualitatively different way as predicted by the deviance model. One strategy looks for the reduced regularity effects that should be apparent in dyslexic children if they are making no use, or less efficient use, of sound-to-spelling or spelling-to-sound rules. The second strategy looks for the deficit in non-word processing that would be expected in dyslexic populations under the same hypothesis. These different experimental strategies have produced conflicting results, for in the case of both reading and spelling most (although not all) recent studies have found no reduction in regularity effects in dyslexic populations when matched with appropriate control groups. Most of these studies have looked at reading (see Brown & Watson, 1991, for a review) and our own experiment described above has found the same pattern for spelling. These results have been interpreted as evidence against a selective deficit in alphabetic processing in dyslexia. However, a much higher proportion of studies have been successful in finding selective non-word reading and spelling deficits in dyslexia, and this is consistent with a wide range of evidence supporting the presence of a general phonological processing deficit in dyslexia. These results do point to qualitatively different processing in dyslexia, supporting a "deviance" model. We have argued that connectionist models of reading and spelling reproduce this apparently paradoxical pattern of effects, and we have suggested that this is because non-word processing is a more sensitive measure of the generalisation performance of a system than the magnitude of regularity effects.

Dual-route and "stage" models

In the case of reading, connectionist models have frequently been interpreted as evidence against "dual route" models of reading, in which there are both lexical and non-lexical routines for synthesising the pronunciation of words. The question arises, therefore, of whether connectionist models of spelling, of the type we have described above, can be seen as potential replacements for traditional rule-based information processing models (mainly dual-route models, in this case). To do this, we now assess the ability of the model to account for the evidence hitherto interpreted as arguing for dual-route models.

Ability to spell irregular words

The fact that people can correctly spell irregular items such as SOAP has been taken as evidence that a single-route model cannot work. However our connectionist model, even though it does not contain two distinct components, can nevertheless learn to spell words with irregular spellings. Further evidence is provided by another connectionist model of spelling development, that of Olson and Caramazza (1988) which also learns to spell both regular and irregular words.

Developmental evidence

The nature of children's spelling errors changes over time. According to the standard "stage" accounts, children go through an initial logographic stage in which they omit letters, and may then spell syllables by the letter whose name is that syllable. At this stage they do not show sound-to-spelling effects. Such regularity effects will emerge in the next stage of spelling, however, as the child begins to grasp the alphabetic principle. Children may incorrectly spell words that they previously spelled correctly as they make the transition between different stages of development. Regularity effects will finally become smaller, and perhaps disappear altogether, as children cease to rely solely on sound-to-spelling translation and acquire knowledge of exceptional spellings.

We have not conducted a detailed analysis of the nature of the errors made by our model, for several reasons. Firstly, the limited vocabulary of the model precludes the meaningfulness of such analysis. Secondly, the nature of the input/output representations that a model uses will be crucial to the particular errors that are produced, and we claim no particular psychological plausibility for the nature of the representations we have chosen. Indeed, it is clear that a complete model would need to use some other representational scheme. However, we have argued (as have Seidenberg & McClelland, 1989a) that the precise nature of the input/output representations is not crucial provided they embody enough of the structure of the input and output domains to enable the model to capture some of the co-occurrence relationships between the two.

Evidence from acquired dysgraphia

One traditional source of evidence for the existence of two separate spelling routines has been the pattern of impairments suffered by brain-injured patients. Some patients ("phonological dysgraphics") selectively 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) involves a relative preservation of the sound-to-spelling translation routine, allowing spelling of regular words and non-words, 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.

Elsewhere (Loosemore, Brown & Watson, 1991) we have shown that “lesioning” the model, after it has learned, can lead to a selective deficit in spelling irregular words similar to that shown by surface dysgraphics (cf. also Olson & Caramazza, 1988). Patterson et al. (1989) have provided a similar demonstration in the case of the connectionist model of reading described earlier. However, it remains to be shown that a unitary connectionist model can handle the complementary pattern of impairment as observed in phonological dysgraphia.

Development of phonemic awareness

In this section we discuss ways in which the connectionist approach could improve our understanding of the development of phonemic awareness. Longitudinal studies (e.g. Cataldo & Ellis, 1988; Mommers, 1987; Shanahan & Lomax, 1986) have demonstrated the importance of spelling in the development of levels of phonemic awareness. One of the key findings of the past decade is that phonemic awareness, first implicit and then explicit, is an excellent predictor of subsequent reading ability. The large literature cannot be more than touched on here, but for present purposes we simply offer some suggestions as to how a connectionist approach might enable us to understand how a computational system could develop phonemic awareness partly as a process of learning to spell.

It is well known that simple connectionist learning algorithms can lead to the establishment of “interesting” and economical representations over layers of hidden units. This applies particularly to sequential networks, which can deal with temporal information flow (Elman, 1988). Indeed, Hanson & Burr (1990) have argued that this integration of ‘learning’ and ‘representation’ is the major contribution of connectionism. It may be that the development of such representations is related to the ability of the system to exhibit implicit “phonemic awareness." These hidden-unit representations can then be recruited to enable spelling development, which will in turn influence the hidden-unit representations and explicit phonemic awareness.

The nature of the representations that are formed will depend upon the task requirements. We argue that spelling imposes different requirements to the prior development of spoken language. More specifically, the temporal segmentation that is required for the former is more specific, or temporally fine-grained. We are currently exploring the possibility that the provision of alphabetic knowledge (in the form of a set of output units representing letters), and the additional requirement to learn sound-spelling mappings (over and above the mappings needed to learn to produce a sequence of phonetic features) can force the development of the more specific phonemic representations over hidden units. The hidden unit representations, which can come to encode temporal features, are then taken as the input to letter-representing output units, and the system is required to learn (using the standard backpropagation algorithm) to produce the correct letter sequences as well as the correct sequences of articulatory features. Thus the ability to learn sound-letter mappings will depend on the specificity of the hidden-unit representations available at any given point in time, and the need to develop spelling ability can in turn motivate the development of "sharper" hidden representations which are comparable to the development of explicit phonemic awareness.

Conclusion

In this chapter we have attempted to show how the use of computational modelling techniques within a connectionist framework can increase our understanding of the development of reading and spelling in normal and developmentally dyslexic children.

The work of Seidenberg and McClelland, and our own connectionist model of spelling reported here and elsewhere, has, we argue, demonstrated a number of points. First, connectionist models can learn to read and spell both regular and irregular words. During learning, the relative difficulty that the models experience with the different word types closely mirrors the level of difficulty experienced on the same words by children. We take this to show that the process of learning to read and spell can usefully be viewed as a statistical one, involving the gradual mastery of associations between patterns in one domain and patterns in another. Further evidence for this conclusion comes from studies of "lesioned" networks - removing computational processing capacity from a system which has already learned a mapping leads to deficits which are qualitatively similar to those experienced by certain brain-injured patients.

Furthermore, we have shown that in conducting experiments looking at sound-to-spelling "regularity" effects it is important to control for the number of "friends" a words has as well as its number of "enemies." Previous experimental studies of spelling regularity have generally confounded these two factors.

With regard to the cognitive processing deficit in developmental dyslexia, we have argued that much of the pattern of difficulty experienced by dyslexics in spelling can be explained in terms of the dyslexic children allocating fewer processing resources to the learning process. Furthermore, the model offers an explanation of a paradoxical finding in the literature - the observation that dyslexics, when matched with appropriate controls, seem more likely to exhibit a selective deficit in non-word processing than a reduced regularity effect, even though both of these are predicted by the hypothesis that dyslexics suffer a selective difficulty in alphabetic processing.

Acknowledgements

Address reprint requests to G. D. A. Brown, Cognitive Neurocomputation Unit, Department of Psychology, University of Wales, Bangor, Gwynedd, LL57 2DG, U.K. The research reported in this paper was supported by grants to the first author from the Medical Research Council (U.K) and the Leverhulme Trust. We are grateful to the schools that allowed us to collect the data reported here, and to Jo Romney, Pippa Glover and Frances Watson for data collection and analysis.

Bibliography

BADDELEY, A.D., LOGIE, R.H. & ELLIS, N.C. (1988). Characteristics of developmental dyslexia. *Cognition*, 29, 197-228.

BARRON, R.W. (1980). Visual and phonological strategies in reading and spelling. In U. Frith (Ed.), *Cognitive processes in spelling*. London: Academic Press.

BARRY, C. & 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.

BEECH, J.R. & HARDING, L.M. (1984). Phonemic processing and the poor reader from a developmental lag viewpoint. *Reading Research Quarterly,* 19, 357-366.

BEAUVOIS, M-F, & DEROUSNE, J. (1981). Lexical or orthographic dysgraphia. *Brain*, 104, 21-50.

BESNER, D., TWILLEY, L., McCANN, R.S. & SEERGOBIN, K. (1990). On the association between connectionism and data: Are a few words necessary? *Psychological Review,* 97, 432-446.

BRADLEY, L., & BRYANT, P. (1981). Visual memory and phonological skills in reading and spelling backwardness. *Psychological Research,* 43, 193-199.

BROWN, G.D.A. (1987). Resolving inconsistency: A computational model of word naming. *Journal of Memory and Language,* 23, 1-23.

BROWN, G.D.A., LOOSEMORE, R., & WATSON, F.L. (1992). Normal and dyslexic spelling: A connectionist approach. Ms. submitted.

BROWN, G.D.A., WATSON, F.L., & LOOSEMORE, R. (1992). Regularity effects in dyslexia: A review, an experiment, and a model. Ms. submitted.

BROWN, G.D.A. & WATSON, F.L. (1991). *Reading in developmental dyslexia: A connectionist approach*. In M. Snowling & M. Thomson, (Eds.), "Dyslexia: Integrating theory and practice" pp. 165-182.

CAMPBELL, R. (1985). When children write non-words to dictation. *Journal of Experimental Child Psychology,* 40, 133-151.

CARAMAZZA, A., MICELI, G., VILLA, G., & ROMANI, C. (1987). The role of the graphemic buffer in spelling: Evidence from a case of acquired dysgraphia. *Cognition,* 26, 59-85.

CATALDO, S. & ELLIS, N.C. (1988). Interactions in the development of spelling, reading and phonological skills. *Journal of Research in Reading,* 11, 86-109.

ELLIS, A.W. (1984). *Reading, writing and dyslexia*. London: LEA.

ELLIS, A. W. & YOUNG, A. W. (1988). *Human cognitive neuropsychology*. London:

LEA.

ELMAN, J.L. (1988). *Finding structure in time*. CRL Technical Report 8801, University of California, San Diego.

FRITH, U. & SNOWLING, M. (1983). Reading for meaning and reading for sound in autistic and dyslexic children. *British Journal of Developmental Psychology,* 1, 329-342.

FRITH, U. (1980). Unexpected spelling problems. In U. Frith (Ed.), *Cognitive processes in spelling*. London: Academic Press.

FRITH, U. (1985). Beneath the surface of developmental dyslexia. In K.E. Patterson, J.C. Marshall & M. Coltheart (Eds.), *Surface dyslexia*. Hillsdale, NJ: LEA.

HANSON, S. J., & BURR D.J. (1990). What connectionist models learn: Learning and representation in connectionist networks. *Behavioral and Brain Sciences*, 13, 471-518.

HATFIELD, F. M. & PATTERSON K. E. (1983). Phonological spelling. *Quarterly Journal of Experimental Psychology,* 35(A), 451-468.

JORM, A.F. (1981). Children with reading and spelling retardation: Functioning of the whole-word and correspondence-rule mechanisms. *Journal of Child Psychology and Psychiatry,* 22, 171-8.

LOOSEMORE, R.L., BROWN, G.D.A., & WATSON, F.L. (1991). A connectionist model of spelling development and developmental and acquired dyslexia. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society,* . Hillsdale, NJ: Lawrence Erlbaum.

McCLELLAND, J. L., & RUMELHART, D. E. (1981). An interactive model of context effects in perception. *Psychological Review,* 88(5), 375-407.

MOMMERS, M.J.C. (1987). An investigation into the relationship between word recognition, reading comprehension and spelling skills in the first two years of primary school. *Journal of Reading Research,* 10, 122-143.

OLSON, A., & CARAMAZZA, A. (1988). *Lesioning a connectionist model of spelling*. Talk given at Venice III: Cognitive neuropsychology and connectionism. October 3-7 1988, Venice, Italy.

PATTERSON, K. E., SEIDENBERG, M. S. & McCLELLAND, J. L. (1989). 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. & PINKER, S. (1988). Wickelphone ambiguity. *Cognition*, 30, 189-190.

RACK, J.P., SNOWLING, M.J. & OLSON, R.K. (1992). The nonword reading deficit in developmental dyslexia: A review *Reading Research Quarterly,* 27(1), 29-51.

RUMELHART, D.E. & McCLELLAND, J.L. (1986). On Learning the Past Tenses of

English Verbs. In J.L. McClelland & D.E. Rumelhart (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2* . Cambridge, MA: MIT Press/Bradford Books.

RUMELHART, D.E., HINTON, G.E. & WILLIAMS, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart and J.L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the microstructure of cognition,Volume 1* Cambridge, Mass.: Bradford Books/ MIT Press.

SCHLAPP, U. & UNDERWOOD, G. (1988). Reading, spelling and two types of irregularity in word recognition. *Journal of Research in Reading,* 11, 120-132.

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

SEIDENBERG, M.S. & 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/ Bradford Books.

SEIDENBERG, M.S. (1989) Visual word recognition and pronunciation: A computational model and its implications. In W. Marslen-Wilson (ed), *Lexical representation and process*. Cambridge, Mass: MIT Press/ Bradford Books.

SEIDENBERG, M.S., BRUCK, M., FORNAROLO, G. & BACKMAN, J. (1985). Word recognition processes of poor and disabled readers: Do they necessarily differ? *Applied Psycholinguistics,* 6, 161-180.

SEIDENBERG, M.S., WATERS, G.S., BARNES M.A. & TANENHAUS, M.K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning and Verbal Behavior,* 23, 383-404.

SEYMOUR, P.H.K. & PORPODAS, C. (1980). Lexical and non-lexical processing of spelling in dyslexia. In U. Frith (Ed.), *Cognitive processes in spelling*. London: Academic Press.

SHALLICE, T. (1981). Phonological agraphia and the lexical route in writing. *Brain*, 104, 413-429.

SHANAHAN, T. & LOMAX, R.G. (1986). An analysis and comparison of theoretical models of the reading-writing relationship. *Journal of Educational Psychology,* 78, 116-123.

SNOWLING, M.J. (1981). Phonemic deficits in developmental dyslexia. *Psychological research,* 43, 219-234.

SNOWLING, M.J., GOULANDRIS, N., BOWLBY, M., & HOWELL, P. (1986). Segmentation and speech perception in relation to reading skill: A developmental analysis. *Cognitive Neuropsychology,* 3, 309-339.

SZESZULSKI, P.A. & MANIS, F.R. (1987). A comparison of word recognition processes in dyslexic and normal readers at two reading-age levels. *Journal of Experimental Child Psychology,* 44, 364-376.

TREIMAN, R. & HIRSH-PASEK, K. (1985). Are there qualitative differences in reading behaviour between dyslexic and normal readers? *Memory & Cognition,* 13, 357-64.

WATERS, G.S. & SEIDENBERG, M.S. (1985). Spelling-to-sound effects in reading: Time course and decision criteria. *Memory & Cognition,* 13, 557-572.

WATSON, F.L. & BROWN, G.D.A. (1992). Single word reading in college dyslexics. *Applied Cognitive Psychology*. 6, 263-272.

Figure captions

Figure One

Regularity effects and lexicality effects in a connectionist model of reading

Figure Two

The error scores to regular and irregular words and non-words for the dyslexic and non-dyslexic models of reading

Figure Three

The error scores to regular and irregular words and non-words for the dyslexic and non-dyslexic models of spelling

Figure Four

Percentage correct responses for regular and irregular words
for dyslexics and controls