Loosemore, R.P.W., Brown, G.D.A. & Watson, F.L. (1991). A neural network model of normal and dyslexic spelling. *Proceedings of International Joint Conference on Neural Networks*, 231-236. Seattle: IEEE Press.

A neural net model of normal and dyslexic spelling†

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

Cognitive Neurocomputation Unit
Department of Psychology, University of Wales
Bangor, Gwynedd LL57 2DG, U.K.
Earn/BITNET: PSSØ3Ø@uk.ac.bangor.vaxa

In this paper we describe a connectionist model of the development of alphabetic spelling, and show that its performance resembles that of dyslexic spellers if its computational resources are restricted during learning. We then describe the results of an experiment which tests the predictions of the model on normal and dyslexic spellers. Our results suggest that dyslexic children show a similar pattern of results to normal children of equivalent spelling age in the spelling of single words. Broad agreement is found between the behaviour of the simulation and the performance of human subjects. The fact that the neural network model provides a good characterisation of the spelling process in dyslexic children suggests that a dyslexic's difficulty can be simply characterised as one of difficulty in mastering the mappings from sound to spelling in English, rather than in terms of qualitatively distinct processing.

Introduction

For psychologists and researchers in other disciplines, neural network models can serve many different roles. In one capacity, they have been used to address findings from neuroscience, by providing models that are consistent with neurobiological constraints. Within computer science, an increased understanding of the computational properties of massively parallel architectures is emerging, and improving our understanding of how co-operative computation can take place in large networks of simple neuron-like computing elements. Within cognitive psychology, building models that can be implemented as computer programs has lead to better-specified and more explicit explanations. Connectionist modelling can also lead to a better understanding of how the statistical properties of a set of data can lead to "friends" and "enemies" effects for individual items, i.e. effects that depend upon the neighbourhood structure of the input patterns. In this paper we describe the implementation of a connectionist model of the development of alphabetic spelling, which was applied to a set of words that differed in terms of the number of "friends" and "enemies" that they had. We were particularly concerned to assess the ability of the model to characterise the spelling of dyslexic, as well as normal, spellers.

---

† 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. We are grateful to the schools which allowed us to collect the data reported here.

Models of spelling development

It has been claimed within psychology that there are two separate routes involved in the derivation of spellings. One route involves the use of sound-to-spelling translation rules, by which regular words can be spelled. A second "lexical" route is required for the correct spelling of irregular words, and to distinguish between non-homographic homophones such as THEIR and THERE. The two-route model of spelling provides a descriptive framework within which many relevant empirical phenomena can be accommodated (see Ellis, 1984, for a review). Some significant evidence comes from the observation of brain-injured patients: a "phonological dysgraphic" may retain the ability to spell both regular and irregular words, while being unable to assign spellings to non-words (Shallice, 1981). This clearly indicates an inability to make use of phonological constraints. "Surface dysgraphics," in contrast, are only able to use phonological information in spelling - they are unable to spell irregular words, and their attempts give rise to errors that are phonic in nature (Beauvois & Derousne, 1981). Frith (1985) provides a sophisticated developmental stage model which charts the way in which spelling and reading interact with one another during development. In this model, spelling plays an important role in the shift from visual/logographic to alphabetic processing (see also Marsh, Friedman, Welch & Desberg, 1981).

Psychological accounts of spelling development often refer to the acquisition of *rules* for sound-to-spelling translation. In particular, it has often been assumed that developmental stages reflect reliance on particular routes (i.e. sets of rules) for sound-to-spelling translation. However, recent work on connectionist modelling seems to show that a network may be able to learn associations between two sets of patterns, such as the orthographic and phonological forms of words, without the need for explicit rules or pathways (Seidenberg & McClelland, 1989a). The present research applies the connectionist approach to the development of alphabetic spelling; we have elsewhere addressed the issue of developmental stages between strategies (Brown, 1990). The model described here deals only with the acquisition of sound-to-spelling knowledge. It does not pass through an initial whole-word spelling stage and it does not have a separate source of lexical information. An initial description of the model as applied to normal spelling is given in Brown, Loosemore, Romney and Watson (1991); here we focus on the characterisation of spelling in developmental dyslexia.

A connectionist model of spelling development

The model was implemented as a three-layer net (with 50 input, 30 hidden and 50 output units) using the backpropagation learning algorithm (Rumelhart, Hinton & Williams, 1986). Patterns presented to the network corresponded to 245 monosyllabic English words, with the pronounced form shown as the input and the spelling shown as the output. Both input and output patterns were constructed using a "triples" scheme similar to that used by Seidenberg and McClelland (1989a) - a word like SOAP (pronunciation /swp/), for example, was represented as the combined activation of the triples *_sw + swp + wp_* (input) and *_SO + SOA + OAP + OP_* (output). Each triple was represented as a randomly-chosen seven-bit pattern over the input or output units - seven input units were turned on for each triple contained in the word. This choice of representation is adequate for this size of model (but cf. Prince & Pinker, 1988).

Predictions of the model

Early studies of spelling tended simply to classify words according to their "regularity" (e.g. Barron, 1980; Seymour & Porpodas, 1980). This classification confuses two separate effects (Brown, 1987), which are governed by the number of "friends" and/or "enemies" that a word possesses. For sound-to-spelling translation, a given word has "friends" if there are other words that are pronounced similarly *and* have a similar spelling (e.g. LOCK has friends DOCK, ROCK, SOCK). Words which are pronounced similarly but spelled differently are called "enemies" (e.g. SOAP has enemies HOPE, COPE, POPE). A word such as LOCK is classified as "regular"; but it differs from the "irregular" word SOAP both in terms of the number of enemies (it has none) and also the number of friends (it has many). A word like BULB, by contrast, has neither friends nor enemies (it has no rhymes at all). We investigated the effects of friends and enemies by including 21 words of each of these three word types in the vocabulary presented to the model - these words included those used in the experiment reported below.

Figure 1(a) shows the model's average "error" score for the three different types of word as it learns. The error score is a measure of the difference between the actual output and the target (correct) output at any given stage in learning. Although the output should be treated with caution owing to the small vocabulary and the non-continuous nature of the possible output errors, the same general pattern of behaviour is observed across a wide range of assumptions. The model learns the spellings of the words with many sound-to-spelling friends most accurately, followed by the words with neither friends nor enemies. The words with only enemies are learned least well. To assess the accuracy of the "spellings" produced by the model, we compared the actual output patterns with patterns corresponding to obvious competitor spellings for individual words. In all cases examined the correct spelling was produced. The same general pattern of behaviour was observed over a wide range of conditions.

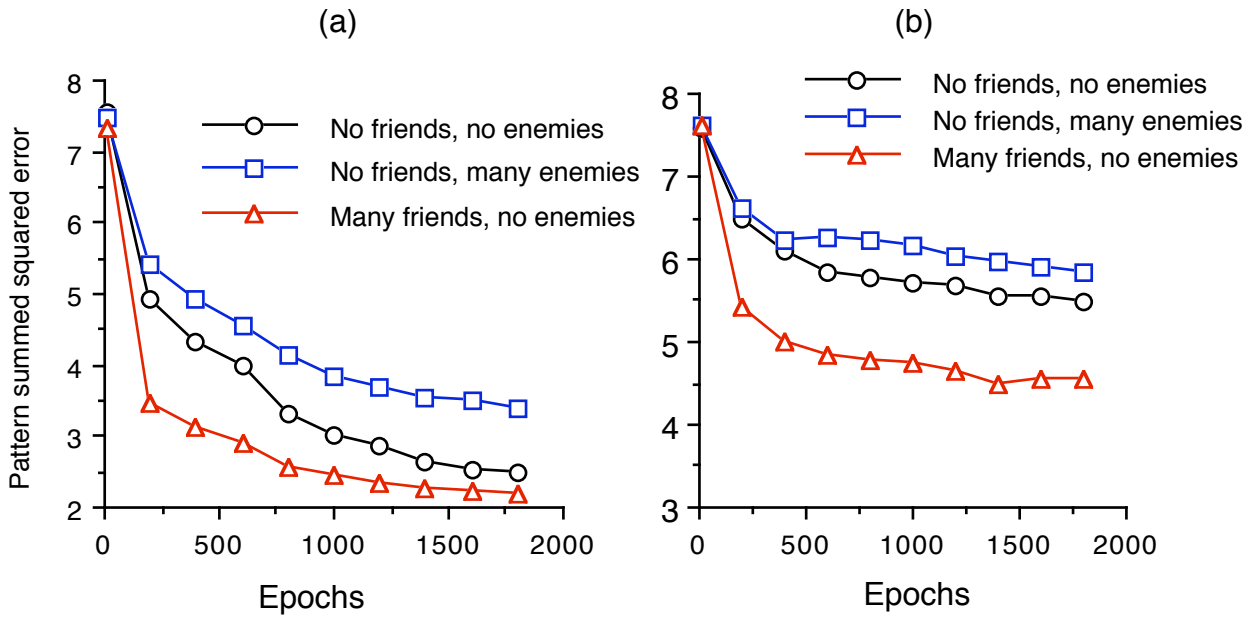

Figure 1. Performance of the model with 30 hidden units (a) and 15 hidden units (b).

The connectionist approach allows us to investigate models of developmental dyslexia by manipulating the available computational resources. In the case of reading, Seidenberg and McClelland (1989b) have suggested that reading backwardness can be characterised in connectionist terms as a reduction in the computational resources - a smaller number of hidden

units - that the network can use to learn the input-output mappings. We therefore examined the effect on our model of providing it with 15 hidden units, rather than 30 as in the previous simulation. The results obtained are shown in Figure 1(b). The basic result is clearly identifiable when the figure is compared with Figure 1(a). It can be seen that the model learns the patterns less well overall with its fewer resources, but that the three types of word are ordered as before. On the assumption that error score may be interpreted as a reflection of human error rate, the model makes a clear prediction about the ordering of word type difficulty in both normal and dyslexic spelling. (An alternative performance measure, taking into account the epoch at which the output pattern becomes criterially close to the target word, is currently under development.) Testing this prediction was the main focus of the experiment described below.

Testing the model's predictions: Normal and dyslexic spelling development

We tested the predictions of the model as applied to normal and dyslexic spelling development. 25 dyslexics from a specialist dyslexic school were tested. Control subjects were chosen from a group of 140 used in an earlier experiment (conducted with Joanna Romney). Dyslexic and control subjects were matched on spelling age; thus the dyslexics were older than the controls. The subjects were drawn from two different classes with different age ranges. There were thus four groups of subjects overall - young ("Junior") dyslexics and controls, and older ("Senior") dyslexics and controls. Within each group, we examined error rates for three different word types - those with friends and no enemies (e.g. TRUCK, cf. LUCK, TUCK, BUCK); those with enemies and no friends (e.g. TYPE, cf. WIPE, RIPE, PIPE), and words with neither friends nor enemies (e.g. BULB). This allowed us to assess friends and enemies effects independently.

The stimulus materials for the experiment comprised 20 matched word triplets. Words within each triplet varied in their sound-to-spelling correspondences but were matched as closely as possible on Kucera-Francis (1967) word frequency, bigram frequency (Solso & Juel, 1980) and on word length. No word in the sample was homophonic with any other English word. 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 conducted a comprehension test and looked at error rates only to words that were known to individual subjects.

Results

The spelling errors made by the four different groups of subjects are shown in Figure 2. There were significant independent effects of sound-to-spelling friends and sound-to-spelling enemies, across all the groups of subjects, both dyslexic and control. Thus the results do not support the suggestion that there is any qualitative difference in the spelling strategies adopted by developmental dyslexics when compared with control subjects spelling at the same level.

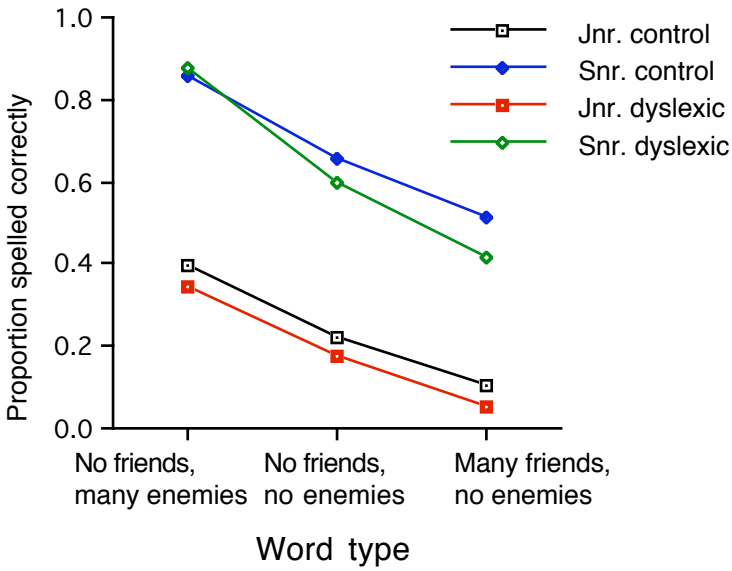

Figure 2.

Proportion of known words spelled correctly by dyslexic and controls.

In both age groups, the performance of the dyslexic children was similar to that of the matched non-dyslexic subjects. Comparison with Figure 1 reveals that our experimental subjects showed a similar pattern of errors to the connectionist model: the version of the model with a restricted hidden unit layer eventually reaches a level of performance that is both quantitatively and qualitatively similar to the performance of the "normal" model at an earlier stage in learning. The "dyslexic" model never reduces its error to the same level as the normal one, which accords with observations of dyslexic humans, who often continue to have trouble spelling some words even into adulthood. The good match between model and observation may be interpreted in terms of the known ability of a backpropagation network to capture statistical regularities in the patterns presented to it - in this case, the English sound-to-spelling mapping system. When the model is required to learn a mapping using limited computational resources, it will choose to represent the regularities that are most strongly present. To the extent that a word has structure in common with others (i.e. it has friends), the encoding of the common regularity will facilitate the learning of that word. Words that are not part of a regular set (having neither friends nor enemies) will not benefit from the representation of the regularities. Words with only enemies will fare worst of all, because the competing orthographic segments associated with the ambiguous phonological segments will have opposite effects on the strengths of connections between units.

Conclusions

The research reported here represents an attempt to provide a computationally explicit account of one aspect of human spelling development. It is intended to supplement rather than to replace more traditional cognitive psychological accounts, by providing a more explicit model which preserves the descriptive-level insights furnished by developmental psychologists. Future models will need to combine both reading and spelling abilities if they are to capture the relevant phenomena from developmental cognitive psychology. However, our implemented connectionist model of spelling has led to a more precise characterisation of sound-to-spelling effects in lexical processing. Our experiment tested the hypothesis that the spelling deficit in developmental dyslexia could be well characterised within a neural network framework in terms of the availability of limited computational resources in the model. The results of the experiment confirm the utility of this approach: developmental dyslexics, when compared to control

subjects spelling at the same level, show identical patterns of performance on the processing of words of several different types.

The connectionist model shows a similar pattern of difficulty across different word types to both normal and dyslexic spellers. The model is simply operating on the basis of the statistical regularities and sub-regularities implicit in the English sound-to-spelling mapping system. In combination, these facts strongly suggest that the process of learning to spell is essentially a statistical one. This perspective contrasts with the more traditional rule-based approach within developmental cognitive psychology.

References

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

BEAUVOIS, M-F, & DEROUSNE, 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, 1-23.

BROWN, G.D.A. (1990). A connectionist model of spelling development. Proceedings of the 1990 International Neural Networks Conference, (Paris, July 1990).

BROWN, G.D.A., LOOSEMORE, R.P.W., ROMNEY, J.L. & WATSON, F.L. (1991). A neural net model of spelling development. Submitted to the International Conference on Artificial Neural Nets, ICANN-91. Helsinki, 1991.

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

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

KUCERA, H., & FRANCIS, W.N. (1967). *A computational analysis of present-day American English*. Providence, RI: Brown University Press.

MARSH, G., FRIEDMAN, M.P., WELCH, V. & DESBERG, P. (1981). A cognitive-developmental approach to reading acquisition. In T.G. Waller & G.E. MacKinnon (Eds.), *Reading research. Advances in theory and practice. Vol 3*. New York: Academic Press.

PRINCE, A. & PINKER, S. (1988). Wickelphone ambiguity. *Cognition*, 30, 189-190.

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.

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.

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.

SOLSO, R.L. & JUEL, C.L. (1980). Positional frequency and versatility of bigrams for two- through nine- letter English words. *Behavior Research Methods and Instrumentation*, 12, 297-343.