Brown, G.D.A., Loosemore, R.P.W., Romney, J. & Watson, F.L. (1991). A neural network model of spelling development. In T. Kohonen, K. Makisara, O. Simula & J. Kangas (Eds.), A*rtificial Neural Networks*. 1727-1730. Amsterdam: Elsevier Science Publishers.

A NEURAL NET MODEL OF SPELLING DEVELOPMENT†

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

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

In this paper we use a simple neural net to model the development of spelling in children, and we describe the results of an experiment which shows that the error rates shown by children in spelling different types of word closely mirror the error scores of the model after it has learned the same words. We conclude that connectionist modelling can inform research into spelling development by focussing attention on the statistical nature of the input-output mappings that children must learn.

Introduction

Within both neural network theory and cognitive psychology, much attention has focussed on the process of translating spelling to sound [1, 2, 3]. This work provides a good example of a fruitful interaction between experimentation and connectionist computational modelling. However, at a computational level, the English sound-to-spelling system provides an equally interesting mapping. Word pronunciations provide only partial constraints on their correct written forms: the spelling of many English words is not entirely predictable from their phonology, because of the ubiquity of words such as SOAP (cf. ROPE, HOPE, COPE, POPE etc). Pairs or triples of words that are pronounced identically but spelled differently pose particular problems for processing (e.g. THEIR - THERE; HARE - HAIR). Although pronunciation does not fully determine orthography, it nevertheless provides useful cues regarding spelling, and phonological information is heavily implicated in human spelling performance. This has led many researchers to postulate models in which two different sets of constraints, reflecting lexical and pronunciation information, are used in determining a word's spelling (see [4] for a review).

Psychological approaches to reading typically make reference to *rules* for translating written representations into phonological codes. However, recent connectionist models can associate different representations (such as orthographic and phonological strings) without reference to explicit rules of any kind. Such models can exploit regularities at many different levels in the mapping from sound to spelling or spelling to sound without making use of rules.
Connectionist modelling can also force a re-consideration of how the global statistical properties of a set of input data can lead to "enemy" and "friend" effects for individual members of the input. Early psychological studies of sound-spelling correspondences tended simply to classify items as

† The research reported in this paper was supported by grants to the first 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.

"regular" or "irregular." This classification is the natural one for accounts framed in terms of the putative sound-to-spelling "rules" of English. However, an alternative explanation for these apparent effects of *regularity* may be given in terms of sound-spelling *friends* and *enemies.* A word like SOAP has an irregular sound-spelling correspondence, because there are other enemy words pronounced similarly but spelled differently (e.g. HOPE, ROPE). However, SOAP is friendless - it is the only -OAP word pronounced and spelled this way. A regular word such as LOCK (cf. DOCK, ROCK) differs from SOAP not only in terms of sound-spelling enemies (it has none) but also in terms of its sound-spelling friends (it has many). This distinction between SOAP and LOCK can be clearly seen by considering a word like BULB, which has neither enemies nor friends. BULB has no rhymes at all so there are no phonological neighbours spelled either the same or differently. In summary, there are words with sound-spelling friends but no enemies (e.g. LOCK), words with enemies but no friends (e.g. SOAP) and words with no friends or enemies (e.g. BULB).

Within psychology, most investigations of spelling have confounded sound-to-spelling friends and enemies, and it is therefore not clear which is the relevant factor in determining performance, although some recent experiments on spelling have found influences of factors other than simple "regularity" on spelling. We suggest that the confusion within the psychological literature is ultimately due to the lack of a sufficient understanding of the statistical nature of the mapping between sound and spelling in English. Accordingly, we used a simple neural network in an attempt to model the development of English spelling knowledge.

A connectionist model of spelling development

The basic architecture of the model involved three layers, with 50 input units, 30 hidden units, and 50 output units. The inputs to the model were representations of the pronunciations of 245 monosyllabic English words. The words included those used in the experiment described below, where the selection of items is described in more detail. Words included those with sound-spelling friends and no enemies (e.g. TRUCK); those with enemies and no friends (e.g. TYPE), and words with neither friends nor enemies (e.g. BULB). This allowed us to assess both friends and enemies effects independently.

The input and output patterns, for the present implementation, were made by superposing patterns for the "triples" that make up a word - SOAP, for example, would be analysed as _SO + SOA + OAP + OP_ (orthographic output) and *_sw* + *swp* + *wp_* (phonological input). The pattern for a triple consisted of four randomly-chosen units ON, and the rest OFF. This type of representation is not without its problems [5] but it suffices for the current model's small vocabulary.

The network learned to associate an input pattern representing the phonology of each word with an output pattern representing its orthography. Backpropagation was used to train the network. 245 words were presented during each epoch, and the net was able to learn the spellings to a reasonable degree of accuracy. As well as calculating error scores against the intended target, comparisons were made against competing spellings of the words (e.g. CLIME is a competitor for CLIMB). In all cases examined, the correct spelling dominated after a few hundred epochs.
Figure 1 shows the model's average pattern summed squared error for three different word types as it learns. The graph represents the average of the error score for all the words in each category. The error is a measure of the difference between the actual output and the target (correct) output at any stage in learning.

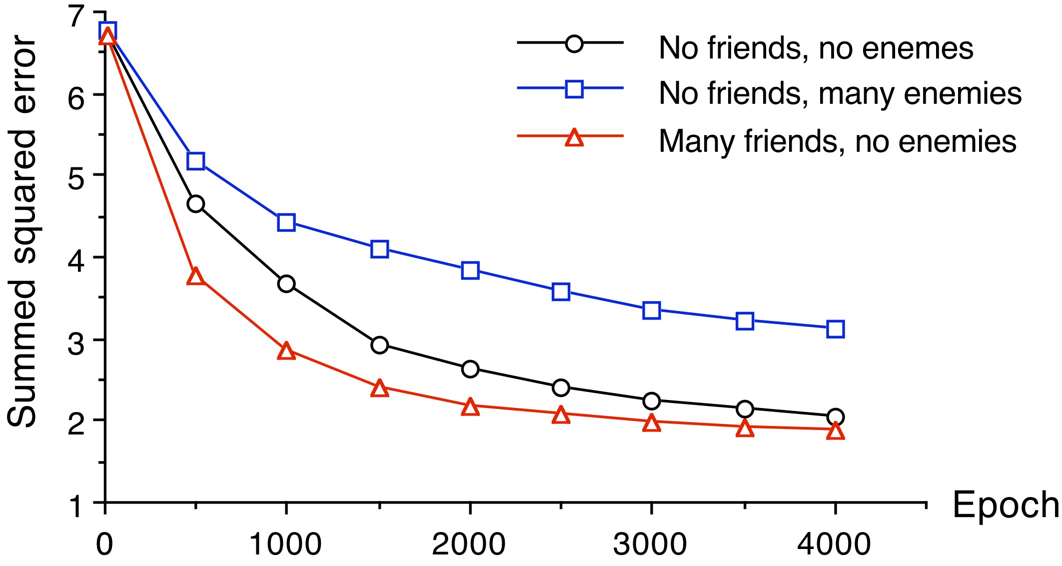

Figure 1: Performance of the model on three word types.

The same general pattern of behaviour is observed across a wide range of assumptions. The model learns the spellings of words with many friends and no enemies best, followed by words with neither friends nor enemies, and words with only enemies are learned least well.

Testing the model's predictions: Normal spelling development

On the assumption that error score in the model may be interpreted as a reflection of human error rate, the connectionist developmental model makes a clear prediction about the ordering of word type difficulty in human spelling. If the relative difficulty experienced by children in spelling different word types reflects the difficulty of learning the statistical mapping between English sound and spelling, then children should show a similar pattern of errors to the model.

We therefore tested the predictions of the model as applied to normal spelling development. We examined the spelling performance of 135 children from four different spelling ability groups, with group mean chronological ages varying from 9:4 years to 10:10 years. 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. CLIMB, cf. SLIME, GRIME, LIME), and words with neither friends nor enemies (e.g. BULB). 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 word frequency, positional bigram frequency, and 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, and the subjects were required to spell the word. We conducted a comprehension test and looked at error rates only to words that subjects knew. The results are presented in Figure 2.

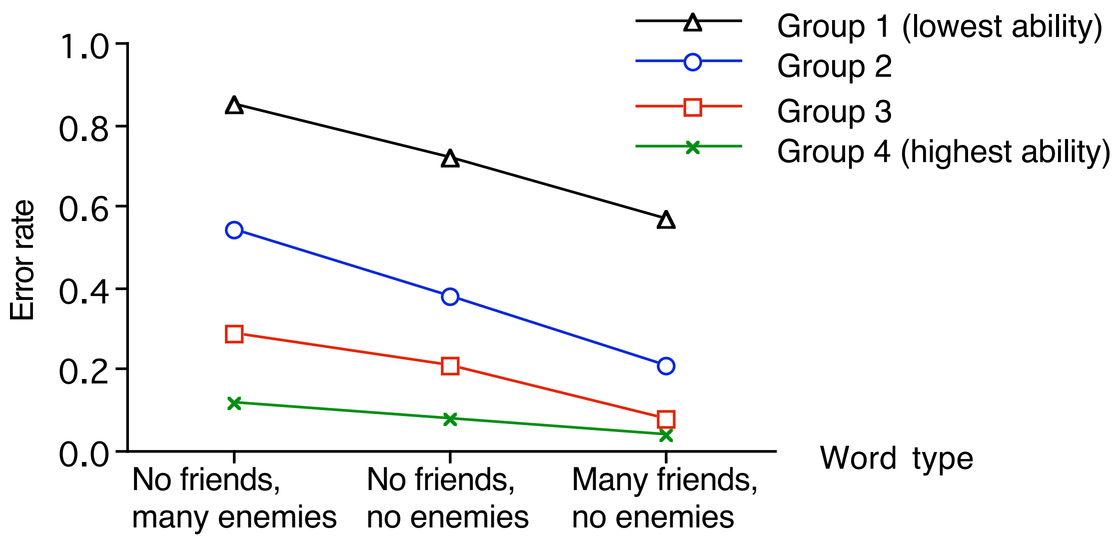

Figure 2: Spelling error rates on three word types by children of four ability levels.

The ordering of word types in terms of errors on known words is similar for the connectionist model and for our subjects. Both showed the highest error rate for words with enemies and no friends, and the lowest error rate for items with friends and no enemies. Both friends and enemies effects were statistically significant.

The good match between model and empirical observation may be interpreted as follows. It is assumed that the behaviour of the connectionist model is simply a reflection of the statistical patterns implicit in the English sound-to-spelling mapping system. When any system with limited computational resources, such as this model, has to learn a set of associations, it will first learn the regularities that are most strongly present in the mapping. Words that are not part of a regular set (those with neither friends nor enemies) will not benefit from the representation of the regularities. Words with enemies will fare even worse, because the competing orthographic segments that are associated with the ambiguous phonological segments will have opposite effects on the strengths of connections between nodes. The similarity between the performance of the model and our human subjects suggests that the behaviour of children as they learn to spell alphabetically can be well characterised as one of mastering a statistical mapping, and that the difficulty of the task will be governed by the structure present in the environment.

References

[1] Seidenberg, M.S., & McClelland, J.L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review,* 96, 523-568.
[2] Sejnowski, T. J., & Rosenberg, C.R. (1986). *NETtalk:* A parallel network that learns to read aloud. Technical Report JHU/EECS-86/1, John Hopkins University.
[3] Brown, G.D.A. (1987). Resolving inconsistency: A computational model of word naming. *Journal of Memory and Language,* 23, 1-23.
[4] Ellis, A.W. (1984). *Reading, writing and dyslexia.* London: LEA.
[5] Prince, A. & Pinker, S. (1988). Wickelphone ambiguity. *Cognition*, 30, 189-190.