



Phonological and Surface Dyslexia in a Single PDP Model of Reading

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Welbourne, SR., & Lambon, R. MA. (2006). Phonological and Surface Dyslexia in a Single PDP Model of Reading. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006, Vancouver)*

Published in:

Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006, Vancouver)

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Phonological and Surface Dyslexia in a Single PDP Model of Reading

Stephen R. Welbourne (stephen.welbourne@manchester.ac.uk)

School of Psychological Sciences, University of Manchester, M13 9PL

Matthew A. Lambon Ralph (matt.lambon-ralph@manchester.ac.uk)

School of Psychological Sciences, University of Manchester, M13 9PL

Abstract

Two simulations were conducted using a replication of PMSP96 (Plaut, McClelland, Seidenberg, & Patterson, 1996, Simulation 4). The first simulation demonstrated that this implementation of PMSP96, was able to reproduce the standard effects of reading, and that when damaged by removal of the semantic input to phonology it produced the kind of frequency/consistency interactions and regularisation errors typical of surface dyslexia. The second simulation explored the effect of phonological damage followed by a period of recovery. This produced large lexicality effects characteristic of phonological dyslexia. This is the first time that symptoms of both of these reading disorders have been demonstrated by the same implementation of the triangle reading model.

Introduction

Phonological dyslexia is a disorder of reading characterised by impairment in nonword reading ability. The characteristics of phonological dyslexia are closely related to those of deep dyslexia with the important distinction that phonological dyslexics do not make any of the semantic errors that are diagnostic of deep dyslexia. The first case of phonological dyslexia was reported by Beauvois and Derouesné (1979) who coined the term. Since then there have been numerous reports of individual cases as well as two case series (Berndt, Haendiges, Mitchum, & Wayland, 1996; Crisp & Lambon Ralph, in press). Analysis of these shows that there is a wide continuum of reading performance both for words and nonwords. At one end there are patients whose word reading is near ceiling and have only slightly impaired nonword reading; then there are patients with relatively 'pure' deficits whose word reading is still reasonably preserved, but whose nonword reading is almost at floor. Finally, there are the very severe cases whose nonword reading is abolished, but who also have poor reading of words.

At first it was thought that the only factor that was important for reading performance in phonological dyslexia was lexicality. More recently, it has been shown that imageability/concreteness also affects word reading. Traditionally this variable has been associated with reading performance in deep dyslexics and most of the early reports do not associate imageability effects with phonological dyslexia. The first suggestion of this possible association comes from patient LB (Derouesné & Beauvois, 1985); however, it was not until the most recent case series study (Crisp & Lambon Ralph, in press) that it became clear that the occurrence of imageability effects in phonological

dyslexia was widespread. In that study all except one of the 12 patients (the mildest) were significantly more accurate when reading high imageability words. This gradual appearance of 'deep dyslexic' symptoms in cases of phonological dyslexia is part of a trend in which deep and phonological dyslexia are viewed as points on a continuum rather than as separate disorders (Friedman, 1996).

Much of the previous work on models of reading has focussed on modelling surface rather than phonological dyslexia (Patterson, Seidenberg, & McClelland, 1989; Plaut et al., 1996); as yet there has been no satisfactory account of acquired phonological dyslexia within a connectionist framework. Harm and Seidenberg (1999) have explored the phenomenon of developmental phonological dyslexia with some success. They trained a single route network in two stages. First they trained the phonological portion of the network so that it learned the phonological representations of the words in the training corpus. They then trained the network to read, interleaving this new training with continued exposure to phonological only trials from the first phase of training. To model developmental phonological dyslexia they damaged the phonological portion of the network after the first stage of training. Although they successfully modelled varying severities of developmental dyslexia, none of their simulations come near to producing the very large lexicality effects found in cases of pure acquired phonological dyslexia. In fact there are no reported PDP models of acquired phonological dyslexia. In view of the inevitable absence of null results in the literature it is difficult to come to any definite conclusion as to why this should be, but we suspect that a key factor is the difficulty in obtaining large lexicality effects. Attempting to model large performance dissociations as a result of damage to a PDP network can be a very frustrating task. Damage to these networks tends to affect all processing tasks with a similar severity. This was certainly the case with early attempts to model surface dyslexia (Patterson et al., 1989): PMSP96 was successful in modelling surface dyslexia, but it achieved this by circumventing the problem. It modelled semantic contributions by applying an external input to 'push' the output of the phonological units towards their targets. Semantic damage could then be modelled by the removal of this input.

This paper adopts an alternative approach to modelling the effects of brain damage (Welbourne & Lambon Ralph, 2005). Under this approach human performance after damage is assumed to be the result of a combination of damage and plasticity-related recovery. The period of recovery (corresponding to the period of spontaneous recovery in patients) is critical because it allows the brain to

re-optimize its remaining connections thus allowing the model to make the best use of what resources it has left. The theoretical position behind that paper is held in common with this study and revolves around the proposition that recovery after brain damage may be, at least in part, attributable to synaptic weight changes. If the human brain's ability to perform accurately depends on the pattern of synaptic weights then it seems reasonable to assume that the removal of a proportion of those weights will not leave the remaining synapses optimally configured to perform the task. Further, provided that there exists some optimisation process by which the synaptic weights can change (learning) then it seems inevitable that some of the recovery that we observe in patients after brain damage must be attributable to synaptic change. This kind of mature synaptic plasticity has been studied mostly in the context of cortical sensory maps (for a review see Buonomano & Merzenich, 1998) and it is clear that these maps are capable of undergoing extensive modification presumably as a result of some learning process operating at the synaptic level.

It seems possible that application of this new methodology to a suitable model might result in a closer match to the symptoms of phonological dyslexia than has hitherto been achieved. We selected Simulation 4 from Plaut, McClelland, Seidenberg and Patterson (1996) as the most appropriate for our purposes. This model consists of a feedforward network trained on a set of monosyllabic words with the training weighted by the square root of word frequency. Input to the phonological units came partly from this network and partly from an external 'semantic' contribution. In their paper, Plaut et al. demonstrated that removal of this semantic contribution resulted in typical surface dyslexic reading patterns. We speculated that damage to the phonological side of the network followed by a suitable period of retraining might result in typical phonological patterns of impairment.

Simulation 1

The architecture, training and representations used in this simulation were modelled on those used by Plaut et al. (Simulation 4, 1996)¹. Each of these key features is summarized below. Figure 1 shows the architecture of the network that was used throughout this set of simulations. There were three sets of units: 105 grapheme units; 100 hidden units and 61 phoneme units. The input layer was connected to the hidden layer with a probability of 40% and the hidden layer was connected to the output layer with a probability of 80%. This sparse connection is a modification from the original simulation where every layer was fully connected to the next layer up. The purpose of this modification was to reduce the competence of the phonological part of the model so that word reading would require a division of labour between semantic and phonological systems. Plaut et al. achieved the same result by using a very high value of weight decay in the phonological part of the model. This method was chosen in preference because it is a more realistic description of

synaptic connectivity in the human brain where connection density is relatively low and dependent on distance (Plaut, 2002; Young, Scannell, & Burns, 1995).

The activity level of each unit was set to vary between 0 and 1 as a nonlinear (logistic) function of the unit's total input. The initial weights on the connections were set to random values between -0.1 and +0.1. The network was then trained using the standard backpropagation learning algorithm with momentum enabled only if the gradient of the error slope was less than 1. Cross entropy was used as the error function as in PMSP96. The learning rate for the network was set to 0.05 and the momentum was 0.9.




Figure 1: Network Architecture.

Orthographic and Phonological Representations

The network used the same representations as PMSP96. These representations divide each word into three parts (onset, vowel and coda) and then use specific units to code for particular graphemes or phonemes occurring within each part.

Imageability Ratings Imageability ratings for words in the corpus were obtained from the MRC Psycholinguistic database and from Cortese and Fugett (2004). Between them these sources provided ratings for 2719 of the 2998 words in the corpus (1529 words had ratings from both sources). For the purposes of this study both these ratings were converted to z scores and averaged if necessary. Words without an imageability rating were given an average imageability value (z score = 0).

Semantic Input Semantic input to the phonological units was provided such that it tended to push the phonological units towards the correct activations. Throughout training the strength of this contribution was gradually increased to mimic the effect of learning. The strength of this input at any given developmental stage was modulated by word frequency and imageability according to the following formula:-

¹ We are grateful to David Plaut for sharing his training patterns with us.

$$IN_s = \left(\frac{0.5}{1 + e^{-(1.14 \text{Log}(\text{Freq}+2)-1)}} + \frac{0.5}{1 + e^{-(\text{Image}_z+1.5)}} \right) \times \text{Mod}$$

Equation 1. Formula for calculating the semantic input to phonological units.

Frequency was taken from Kuçera and Francis (1967) and imageability z score was calculated as above. (The constants in this formula were selected to provide a sensible distribution across the frequency and imageability values in the corpus with more of the variation originating in imageability.) Over the course of development the total semantic input was modulated by an epoch dependent modulation factor that varied from 0.6 to 4.8 in steps of 0.6 where a step occurred after every 200 epochs.

In the case of the nonwords Plaut et al. did not provide any semantic contribution. This may not have been the correct choice for the following reason: in the brain the connections between O and S (either direct O→S or indirect O→P→S) cannot be selectively turned off for nonwords. Hence nonwords will generate some kind of activation across the semantic units which will, in turn, contribute to the activation of phonological units. This nonword semantic activation will not correspond to any known semantic targets (except in the case of lexicalization errors); rather it will represent some kind of average semantic activation for all the visually similar words. This will result in a contribution from semantics to phonology that is effectively random noise. Accordingly, for nonword reading, semantic input was randomly added to the phonological units where the input for each unit varied between -0.5 and +0.5 modulated by the same modulation function as for the real words.

Training Procedure

The network was trained using full batches with the same corpus of 2998 monosyllabic words used in PMSP96. The root frequency (Kuçera & Francis, 1967) of each word was used to scale the error derivatives for the purposes of backpropagation. This has the same effect as using frequency to determine the probability of a word being presented for training; however, it has the considerable advantage that every word can still be presented once every epoch thus considerably compressing the required training time (See Plaut et al, 1996 for a fuller discussion of this issue). To eliminate the possibility that the results might be a consequence of one particular set of initial weights, the network was trained ten times; each time using a different random set of weights as the starting point. These ten trained networks then formed the starting point for further investigations.

Testing Procedure

Seven sets of test stimuli were used to evaluate the network's performance: high frequency regular; high frequency irregular; low frequency regular; low frequency

irregular; regular nonwords; high imageability words and low imageability words.

The regular and irregular words were taken from Taraban and McClelland (1987) and were matched across groups for frequency. The regular nonwords were taken from Glushko (1979) and were created by changing the onset of an existing regular word. These stimuli are the same as those used in PMSP96 so that it is possible to make a direct comparison of results.

The high and low imageability word sets were constructed for the purposes of this simulation. Low imageability words (imageability rating 200-400) were selected from the training corpus and matched, pairwise, on frequency with high imageability words (500-700) also selected from the corpus.

In addition to the performance on the seven sets of test stimuli, the percentage of regularisation errors made by the network on the two irregular stimuli sets was also recorded.

Initial Training

By the end of training the network correctly pronounced all of the words in its corpus including all of the homographs. This is slightly better than the performance achieved by PMSP96, which was 99.7% accurate in word reading. For nonword reading the model was correctly reading 93.02% of the regular nonwords. This is not as good as the 96.5% achieved by PMSP96, but it is nearer to human performance which averages 93.8% (Glushko, 1979).

It is important to verify that this model could replicate the standard frequency/consistency interaction found in the naming latencies of normal human populations (e.g. Seidenberg, 1985; Seidenberg, Waters, Barnes, & Tanenhaus, 1984). Error scores from the network at epoch 2000 were submitted to a 2 x 2 ANOVA where frequency and consistency were treated as between group variables. This confirmed that there was indeed a significant frequency/consistency interaction ($F(1,1916) = 238.1, p < 0.001$). In addition, there were significant main effects of both frequency ($F(1,1916)=306.4, p<.001$) and consistency ($F(1,1916)=521.9, p<.001$). The nature of the interaction was for frequency to be almost completely modulated by consistency. For irregular words low frequencies resulted in a much higher error scores, but for regular words there was almost no effect of frequency. This is consistent with the standard effect found in human reading latencies and with the results found for PMSP96.

In addition to standard effects of consistency and frequency one might also expect to see an effect of imageability (Strain, Patterson, & Seidenberg, 1995) with high imageability items having lower error scores than low imageability ones. To test this, error scores from the high and low imageability word sets were compared. The mean error score for high imageability items was 0.0082 (SD=0.013) whilst the mean error score for the low imageability items was 0.0223 (SD=0.0416). Submitting these scores to a t test revealed that there was, as predicted, a significant difference ($t=-7.08, df=570, p<0.001$).

Surface Dyslexia – Replication of PMSP

Before investigating the possibility that this model can simulate the symptoms of phonological dyslexia it is important to verify that like PMSP96 it is capable of replicating the symptoms of surface dyslexia. Surface dyslexia is characterised by poor reading of low frequency exception words, coupled with accurate reading of nonwords. Errors made in reading irregular words tend to be regularisations or LARCs (Patterson, Suzuki, Wydell, & Sasanuma, 1995); for example reading PINT to rhyme with MINT. To mimic the effect of semantic damage we gradually reduced the strength of the semantic contribution whilst simultaneously adding random noise to it. This was achieved by decreasing the strength of the modulation factor from 4.8 to 0 in ten steps of 0.48 while simultaneously adding increasing amounts of Gaussian noise with a standard deviation increasing in ten steps of 0.75.

Figure 2 shows the results of this simulation. For clarity the regular high frequency, high imageability and low imageability word sets have been omitted – performance on these word sets is very similar to that for low frequency regular words. Low frequency irregular words are the most affected by this manipulation with performance dropping to 53% for the worst damage. At this point performance on high frequency irregular words is reduced to 89%; regularisation errors constitute 79% of all errors made on irregular words; while accuracy rates on all other word sets fall between 95% and 98%. Note that for nonwords this represents a slight improvement on the undamaged performance. This pattern of results is consistent with that found in surface dyslexic patients and with the results of PMSP96 Simulation 4.

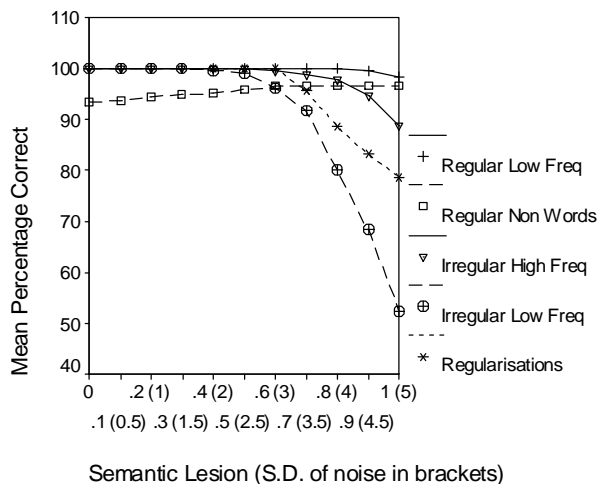


Figure 2. Effect of removal of semantic input

Simulation 2 – Phonological Damage

The architecture, network dynamics and training environment used in this simulation were identical to that of Simulation 1. Starting from the same ten fully trained networks we explored the effect of damage to the phonological portion of the network followed by a period of

retraining. Phonological damage was simulated by lesioning the links between input and hidden layers whilst simultaneously adding noise to the output of the hidden layer. Three levels of damage severity were tested (15%, noise SD=.15; 30%, noise SD=.3 and 70%, noise SD=0.7). After damage the network was allowed to recover for 200 epochs by re-exposing it to the original learning environment. Figure 3 shows the results of this investigation for the three levels of damage severity. At the most severe level nonword reading is abolished while word reading accuracy varies between 20% and 40% depending on the stimuli set. The high imageability and regular high frequency words are read with the highest accuracy while the low imageability and irregular word sets are read with the least accuracy. This pattern of results is what one might expect to see in a rather severe case of phonological dyslexia. The only possible criticism would be that there seems to be an effect of consistency as well as one of imageability; an issue that will be addressed in more detail in the discussion.

For the medium and mild levels of damage the pattern of performance for words is very similar to that for severe damage except that it is centred around progressively higher mean scores: in the case of 30% damage scores range from 57% to 85%, while for milder 15% damage they range from 83% to 97%. In all cases irregular and low imageability words are read less accurately than regular and high imageability words. Nonword reading is seriously impaired for all levels of damage with overall level of nonword reading accuracy decreasing with increasing damage severity. However, even at mild levels of damage nonword reading accuracy is still only 47%.

To confirm the significance of the apparent effects of lexicality, imageability and consistency we submitted the results to a series of t tests: lexicality was tested by comparing performance on low frequency regular words with performance on regular nonwords; imageability was tested by comparing performance on the high and low imageability word sets; consistency was tested by comparing the low frequency regular and irregular word sets. All of these comparisons demonstrated highly significant differences (all p 's < 0.001).




Figure 3. Performance after phonological damage and recovery

Discussion

Two simulations were conducted using a network architecture similar to PMSP96 (Simulation 4). The first simulation demonstrated that our implementation performs similarly to PMSP96, in that it can reproduce the cardinal features of normal reading, as well as the symptoms of surface dyslexia. The second simulation explored the possibility that damage to the phonological portion of the model, followed by a period of recovery would lead to performance resembling that found in phonological dyslexia. This simulation demonstrated that a full range of lexicality effects could be modelled; coupled with the imageability effects that are characteristic of phonological dyslexia. This is the first time that such large lexicality effects have been modelled in a network which also has the capacity to learn. Moreover, it is the first time that simulations of surface and phonological dyslexia have been produced from the same connectionist architecture.

One slightly unexpected aspect of these results is the persistence of a consistency effect following phonological damage. This is not traditionally associated with phonological dyslexia. However, although it is not often reported, phonological dyslexics do often exhibit consistency effects. A re-analysis of data from Berndt et al² (1996) reveals that 9 out of 10 of the patients in the series showed more accurate reading of regular than of irregular words with the performance difference ranging from 2% to 20%. When data from all of the patients are submitted to statistical analysis these differences are shown to be significant ($t=2.32$, $df=9$, $p=0.023$, one tailed). Data from the only other case series of phonological dyslexics (Crisp & Lambon Ralph, in press) is even more emphatic; 10 out of 12 patients showed a superiority for regular words varying from 5% to 33% and the group as a whole showed a

very significant consistency effect ($t=4.41$, $df=11$, $p<0.001$, one tailed). The mean size of the consistency effect for the two sets of patients (including those who did not exhibit a consistency effect) was 5% for the Berndt et al. set and 14% for the Crisp and Lambon Ralph set. This compares with a mean consistency effect of 16% for the network (averaged across all damage severities). In the light of this it seems reasonable to suggest that this simulation has captured a hitherto unremarked feature of phonological dyslexia.

These results pose two important questions: (1) What are the critical components in these simulations that are essential to successfully modelling phonological dyslexia? (2) How do these results mesh with those reported by Welbourne and Lambon Ralph (2005)?

Two features of these simulations seem likely to have significantly contributed to their success in modelling phonological dyslexia. Firstly, the fact that the phonological damage was generalised in nature, affecting both the ability of the network to map from orthography to phonology and the integrity of its phonological representations. This was achieved by combining damage to the connections in the O→P pathway with noise added to the output of the phonological hidden units. Without the addition of noise it is probable that the network would have been able to recover by finding solutions that relied more on the regularities in the training set; resulting in reduced lexicality effects and an increased influence of consistency. The idea that phonological dyslexia arises from generalized phonological damage is consistent with the primary systems hypothesis (Patterson & Lambon Ralph, 1999), which assumes that reading is subserved by the more general pre-existing language systems and that the acquired dyslexias arise from generalized damage to one of these systems. Indeed, the current model could be regarded as a first step towards an implementation of the primary systems hypothesis. Of course, a full implementation would require a model that was able to perform additional linguistic tasks such as speech, comprehension and repetition.

The second key factor in this simulation is the inclusion of a period of recovery after damage. Welbourne and Lambon Ralph (2005) found that including a period of recovery was helpful when modelling surface dyslexia because it magnified the effect of small pre-existing processing biases into large performance dissociations. Exactly the same effect is produced in these simulations, but this time the biases are towards lexicality and imageability effects rather than a frequency/consistency interaction.

It is important to consider how the results of this simulation mesh with the results reported by Welbourne and Lambon Ralph (2005). In that simulation, damage to an isolated phonological network resulted in a surface dyslexic performance; here, on the other hand, surface dyslexia arises from damage to the semantic portion of the network whilst damage to the phonological portion produced the symptoms of phonological dyslexia. At first glance this seems somewhat inconsistent; how is it that surface dyslexia can arise from two different damage loci? In reality, there is no inconsistency; in both cases the endpoint is the same. Surface dyslexia occurs where the phonological system has insufficient computational resources to successfully process

² We are grateful to Rita Berndt for allowing us access to some of the data that formed the basis of the paper.

all of the words in its corpus and has no available support from semantics. Welbourne and Lambon Ralph (2005) achieved this situation by damaging a phonological system that was initially over competent in that it could read without any support from semantics. In the current simulation the same situation was achieved, more realistically, by removing semantics from a network where reading was supported by a division of labour between phonology and semantics. Only in this latter situation, where there is the potential for a division of labour, can damage to the phonological system result in phonological dyslexia.

This study represents a considerable step forward in that it is the first time that an implementation of the triangle model has been able to produce both the frequency/consistency interactions typical of surface dyslexia and the lexicality/imageability effects associated with phonological dyslexia.

Acknowledgments

This research was supported by grants from NIMH (No. P50-MH64445) and BBSRC (S20390).

References

- Beauvois, M. F., & Derouesné, J. (1979). Phonological alexia: three dissociations. *Journal of Neurology, Neurosurgery, and Psychiatry*, *42*, 1115-1124.
- Berndt, R. S., Haendiges, A. N., Mitchum, C. C., & Wayland, S. C. (1996). An Investigation of Nonlexical Reading Impairments. *Cognitive Neuropsychology*, *13*(6), 763-801.
- Buonomano, D. V., & Merzenich, M. M. (1998). Cortical Plasticity: From synapses to maps. *Annual Review of Neuroscience*, *21*, 149-186.
- Cortese, M. J., & Fugett, A. (2004). Imageability ratings for 3,000 monosyllabic words. *Behaviour Research Methods, Instrumentation, & Computers*, *36*(3), 384-387.
- Crisp, J., & Lambon Ralph, M. A. (in press). Unlocking the nature of the phonological-deep dyslexia continuum: the keys to reading aloud are in phonology and semantics. *Journal of Cognitive Neuroscience*.
- Derouesné, J., & Beauvois, M. F. (1985). The "phonemic" stage in the non-lexical reading process: Evidence from a case of phonological alexia. In K. Patterson, J. C. Marshall & M. Coltheart (Eds.), *Surface Dyslexia* (pp. 399-457). Hove, UK: Lawrence Erlbaum Associates Ltd.
- Friedman, R. B. (1996). Recovery from Deep Alexia to Phonological Alexia: Points on a Continuum. *Brain and Language*, *52*, 114-128.
- Glushko, R. J. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, *5*(4), 674-691.
- Harm, M. W., & Seidenberg, M. S. (1999). Phonology, Reading Acquisition, and Dyslexia: Insights From Connectionist Models. *Psychological Review*, *106*(3), 491-528.
- Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Patterson, K., & Lambon Ralph, M. A. (1999). Selective disorders of reading? *Current Opinion in Neurobiology*, *9*(2), 235-239.
- Patterson, K., Seidenberg, M. S., & McClelland, J. L. (1989). Connections and disconnections: Acquired dyslexia in a computational model of reading processes. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neuroscience*. London: Oxford University Press.
- Patterson, K., Suzuki, T., Wydell, T., & Sasanuma, S. (1995). Progressive aphasia and surface alexia in Japanese. *Neurocase*, *1*(2), 155-165.
- Plaut, D. C. (2002). Graded modality-specific specialisation in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, *19*(7), 603-639.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding Normal and Impaired Word Reading: Computational Principles in Quasi-Regular Domains. *Psychological Review*, *103*(1), 56-115.
- Seidenberg, M. S. (1985). The time course of phonological code activation in two writing systems. *Cognition*, *19*, 1-10.
- 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 Behaviour*, *23*(383-404).
- Strain, E., Patterson, K., & Seidenberg, M. S. (1995). Semantic Effects in Single-Word Naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(5), 1140-1154.
- Taraban, R., & McClelland, J. L. (1987). Conspiracy Effects in Word Pronunciation. *Journal of Memory and Language*, *26*(608-631), 608-631.
- Welbourne, S. R., & Lambon Ralph, M. A. (2005). Exploring the impact of plasticity-related recovery after brain damage in a connectionist model of single-word reading. *Cognitive, Affective & Behavioral Neuroscience*, *5*(1), 77-92.
- Young, M. P., Scannell, J. W., & Burns, G. (1995). *The Analysis of Cortical Connectivity*. Heidelberg: Springer-Verlag.