The role of plasticity-related functional reorganization in the explanation of central dyslexias

Stephen R. Welbourne [a] , Anna M. Woollams [a] , Jenni Crisp [b] & Matthew A. Lambon Ralph [a]

[a] School of Psychological Sciences, University of Manchester, Manchester, UK

[b] Sheffield Teaching Hospitals NHS Trust, Sheffield, UK

Available online: 28 Nov 2011

PLEASE SCROLL DOWN FOR ARTICLE

Ψ Psychology Press
Taylor & Francis Group

The role of plasticity-related functional reorganization in the explanation of central dyslexias

Stephen R. Welbourne[1], Anna M. Woollams[1], Jenni Crisp[2], and Matthew A. Lambon Ralph[1]

[1]School of Psychological Sciences, University of Manchester, Manchester, UK
[2]Sheffield Teaching Hospitals NHS Trust, Sheffield, UK

This investigation explored the hypothesis that patterns of acquired dyslexia may reflect, in part, plasticity-driven relearning that dynamically alters the division of labour (DOL) between the direct, orthography → phonology (O → P) pathway and the semantically mediated, orthography → semantics → phonology (O → S → P) pathway. Three simulations were conducted using a variant of the triangle model of reading. The model demonstrated core characteristics of normal reading behaviour in its undamaged state. When damage was followed by reoptimization (mimicking spontaneous recovery), the model reproduced the deficits observed in the central dyslexias—acute phonological damage combined with recovery matched data taken from a series of 12 phonological dyslexic patients—whilst progressive semantic damage interspersed with recovery reproduced data taken from 100 observations of semantic dementia patients. The severely phonologically damaged model also produced symptoms of deep dyslexia (imageability effects, production of semantic and mixed semantic/visual errors). In all cases, the DOL changed significantly in the recovery period, suggesting that postmorbid functional reorganization is important in understanding behaviour in chronic-stage patients.

Keywords: Reading; Connectionist; Plasticity; Dyslexia; Parallel distributed processing.

The idea that the division of labour between the direct (orthography → phonology; O → P) pathway and the semantically mediated (orthography → semantics → phonology; O → S → P) pathway is key to our understanding of reading has long been an important strand in theoretical models. The debate encompasses a range between two extreme positions: that reading is primarily a phonological decoding process; or that reading is primarily a recognition process. The consensus opinion has swung between these two extremes (Coltheart, 1978; Frost, 1998; McCusker, Hillinger, & Bias, 1981; Smith, 1973), although more recent modelling studies

65

(Harm & Seidenberg, 2004; Plaut, McClelland, Seidenberg, & Patterson, 1996) suggest that both processes are required to explain the full spectrum of reading phenomena and that the division of labour between phonology and semantics varies according to the nature of the item being processed.

Previous work has considered division of labour (DOL) in terms of how the balance emerges during normal learning and how this premorbid DOL may affect postmorbid performance (Harm & Seidenberg, 2004; Plaut, 1997; Plaut et al., 1996). What has not previously been considered is the possibility that plasticity-driven postmorbid relearning processes may dynamically alter the DOL during the period of recovery and that this may contribute significantly to the eventual pattern of chronic performance observed in patients with acquired reading disorders. The goal in the present series of simulations was to explore how changes in the division of labour during the recovery phase might help to account for the key symptoms of the central dyslexias (Shallice & Warrington, 1980).

As the purpose of this paper is to explore the computational plausibility of this novel hypothesis, it is not appropriate to make a detailed consideration of alternative approaches (since none of the other approaches considers reorganization). The work we report here follows some general principles common to most other computational work on oral reading, in that it assumes that phonological output is driven by two types of input: one based directly on the learned mappings between orthography and phonology and the other on indirect mediation via semantic representations. In common with other triangle-based approaches, we assume that both of these pathways operate simultaneously using parallel processing throughout.

Given that the current model has been inspired by and implemented many aspects of previous versions of the triangle model, it is clearly important to consider the history of these triangle-based models in some detail. This is done in the following sections, where the models are considered alongside the particular central dyslexia that they aimed to simulate.

Dyslexic reading targets

Shallice and Warrington (1980) first coined the term "central dyslexias" to describe those disorders of reading that arise from impairments either to the language system itself or to the mapping between the visual word form (orthography) and the language system. This includes surface dyslexia, phonological dyslexia, and deep dyslexia, but excludes neglect, visual, attentional, and pure alexia (the peripheral dyslexias). The reading performance observed amongst patients with one of the central dyslexias was the target for this study.

Surface dyslexia
Surface dyslexia is characterized by a strong frequency−consistency[1] interaction in reading accuracy with especially poor performance on low-frequency words with inconsistent spellings.

---

[1] This interaction is sometimes also referred to as frequency−regularity or frequency−typicality. Regularity and consistency are heavily confounded in English but they are theoretically distinct: Regular words follow the standard grapheme-to-phoneme conversion rules for pronunciation, whereas consistent words have word bodies whose pronunciation is consistent with their neighbours. In our stimuli, these properties are indistinguishable, as all the regular stimuli are also consistent. However, in the light of this and the considerable amount of theoretical effort that has been put into distinguishing regularity and consistency effects (Andrews, 1982; Cortese & Simpson, 2000; Jared, 2002; Jared, McRae, & Seidenberg, 1990; Jefferies, Ralph, Jones, Bateman, & Patterson, 2004; Taraban & McClelland, 1987), we have elected to use the term "consistency" to describe the phenomenon, as this reflects our belief that that these phenomena result from graded variation in the orthography-to-phonology mappings found in the language at multiple subword levels, rather than in the application of dichotomous rules at the grapheme−phoneme level. It is worth noting that nonword letter strings can also vary in their consistency (Zevin & Seidenberg, 2006). Similarly, we shall use the term "legitimate alternative reading of the components" (LARC; Patterson et al., 1995) rather than the more traditional term "regularization" to refer to errors where the phonological output when reading inconsistent words is what would be expected from analogy with neighbours.

In addition, there is a marked tendency for the errors made on these words to be LARC (legitimate alternative reading of the components) errors (LARC errors include any legitimate alternative reading of the components—for instance, BLOOD might be read to rhyme with GOOD or FOOD; Patterson, et al., 1995; Woollams, Lambon Ralph, Plaut, & Patterson, 2007). The first modern cases of surface dyslexia were reported by Marshall and Newcombe (1973) and were followed up by well-known single-case studies such as M.P. (Bub, Cancelliere, & Kertesz, 1985), and K.T. (McCarthy & Warrington, 1986).

The largest number of reported cases of surface dyslexia have come from a pool of patients who all suffer from the same type of progressive illness known as semantic dementia (K. S. Graham, Hodges, & Patterson, 1994; N. L. Graham, Patterson, & Hodges, 2000; Jefferies, et al., 2004; McCarthy & Warrington, 1986; Patterson & Hodges, 1992; Patterson, Plaut, Seidenberg, Behrmann, & Hodges, 1996). Neurologically this illness involves gradual atrophy focused on the inferolateral anterior temporal lobes, bilaterally (Mummery et al., 2000). Behaviourally, this results in a range of symptoms, all stemming from the loss of semantic knowledge. These include a semantically driven anomia, poor performance on verbal and nonverbal semantic tests, surface dysgraphia, and surface dyslexia (Hodges & Patterson, 2007; Patterson et al., 2006). As the severity of the atrophy increases, so does the degree of semantic impairment and the resultant language symptoms (Lambon Ralph, McClelland, Patterson, Galton, & Hodges, 2001). In terms of their reading performance, these patients move from mild to severe surface dyslexia as their disease progresses (Patterson & Hodges, 1992). By far the largest survey of reading in this group of patients can be found in work by Woollams et al. (2007). This study consisted of 100 observations of reading in a group of 51 semantic dementia patients. The article demonstrated that there was a very strong correlation between reading performance and accuracy on semantic tasks (naming and comprehension). The strength of this correlation was modified by word type and was greatest for low-frequency words with inconsistent spelling–sound correspondences.

Patterson, Seidenberg, and McClelland (1989) were the first to attempt to model surface dyslexia using a computational framework. They took the "triangle model" developed by Seidenberg and McClelland (1989) and explored the effects of lesions at different locations between orthography and phonology. (It should be noted that although this model was inspired by the triangulate of orthography, O, phonology, P, and semantics, S, it only implemented the O → P portion of the network.) Although there was some match between the performance of the model under damage and that found in surface dyslexia, a number of limitations were also apparent. In particular, the model did not display a sufficiently large frequency–consistency interaction when damaged. Neither did it produce enough LARC errors to make a convincing case that it was modelling surface dyslexia of the kind found in fluent patients, although it was a reasonable simulation of the pattern of reading impairments sometimes found in dysfluent patients, who tend to show smaller frequency–consistency interactions and more generalized word-reading impairments (Seidenberg, 1995).

In subsequent simulations, Plaut et al. (1996) tested the hypothesis that successful modelling of surface dyslexia would require a system with input from semantics—surface dyslexia could then be modelled by the removal of semantic input (Simulation 4). They implemented this by training a network that mapped from O → P in the context of an additional semantic input to phonology (mimicking the effect of semantics). When this input was removed, the network's performance very closely resembled that of surface dyslexia. Varying the degree of semantic damage allowed the network to simulate the performance of classic cases (M.P. and K.T.) who between them span the range of surface dyslexic performance.

Two key insights from this study are that successful modelling of surface dyslexia depends on the network achieving an appropriate division of

labour between the direct and semantic pathways, and that premorbid differences in this DOL can result in different patterns of surface dyslexic performance when semantics is removed. For high-frequency words or those with consistent pronunciations, the bulk of the reading computation can be underpinned and is most efficiently accomplished by the direct O → P pathway—and in these circumstances, the semantic contribution, though present for all real words, is superfluous. Where pronunciations deviate from the consistent pattern, then the additional semantic constraint is especially helpful. Hence within this framework, when semantic input is removed, the result is surface dyslexia. This idea was further developed by Woollams et al. (2007) who used a very similar model to show that it could account for data from 100 observations of surface dyslexic patients. The key modification to the original model was that as the semantic input was reduced, it also had increasing levels of Gaussian noise added to it, providing a more realistic implementation of semantic damage.

Deep dyslexia

The first explicitly identified deep dyslexic case was reported by Marshall and Newcombe (1973) in the same paper as that reporting the first modern surface dyslexic. However, deep dyslexia was described in much more detail in a paper by the same authors (Marshall & Newcombe, 1980) in which they surveyed a number of older reported cases who shared the common association of symptoms that define deep dyslexia (semantic errors, derivational errors, visual errors, poor function word reading, and poor or nonexistent nonword reading). Coltheart (1980) expanded this list of symptoms further, importantly including an imageability effect in the clinical profile. Whilst deep dyslexia invariably includes this collection of features, the production of semantic reading errors is the most striking and diagnostically important of them.

Hinton and Shallice (1991) and later Plaut and Shallice (1993) provided the only reading models that address deep dyslexia in any depth (Farrar & Van Orden, 2001, presented a model that can simulate a single semantic error, but their training corpus only included 13 words, precluding a detailed account of the disorder). Both models used an architecture that linked orthography to semantics via a set of hidden units. Damage to this architecture resulted in error types that are comparable to those found in deep dyslexic patients with the same co-occurrence of semantic and visual errors that is peculiar to deep dyslexia. The key insight from this model was that the characteristic co-occurrence of visual/phonological[2] and semantic errors did not require the assumption of simultaneous damage to the visual/phonological and semantic systems. Instead, both could emerge naturally from damage to a system in which the visual and semantic systems were interactively linked. The Plaut and Shallice model also successfully demonstrated an imageability effect by modelling high-imageability words as having more semantic features than low-imageability words. As far as we are aware, with the exception of the current model, this is the only other computational model that can simulate all of these key deep dyslexic symptoms. Despite its considerable successes, this model is only partially helpful in meeting our current goal to explain all types of central dyslexia simultaneously because the target of the previous work was deep dyslexia alone. As a consequence, it was not necessary to implement a computational architecture that included a direct mapping between orthography and phonology because deep dyslexic patients' nonword reading accuracy is typically at floor. The Plaut and Shallice architecture is silent, therefore, on various core targets for our simulations, including the emergent division of labour (see above), the ability of the model to read nonwords, and thus the capacity to simulate surface and phonological dyslexia.

---

[2] These are described in the original paper as visual errors, but visual and phonological errors are very hard to distinguish in English due to the largely regular nature of the relationship between orthography and phonology. We refer to them throughout as visual/phonological errors.

Phonological dyslexia

The insistence on the central importance of semantic errors in deep dyslexia meant that cases of patients who were similar in many respects to mild deep dyslexics (e.g., patient A.M.; Patterson, 1982) but who did not produce any semantic errors were given a different label. The first such case was reported by Beauvois and Dérouesné who used the label "phonological dyslexia" for this type of patient since they appeared to have a "disturbance of the phonological reading process" (see also Patterson & Marcel, 1992, for evidence that this is primarily a phonological deficit).

More recently, it has been suggested that phonological and deep dyslexia may form a continuum with phonological dyslexia at the mild end and deep dyslexia at the severe end (Glosser & Friedman, 1990). Friedman (1996) studied the recovery profile of five deep dyslexic patients and concluded that there was a strict order of symptoms along a severity continuum. According to this list, semantic errors are associated with the most impaired patients, whereas poor nonword reading would be present in even the mildest cases. In keeping with this proposal, semantic errors were the first symptom to disappear in recovery, whereas, at the other end of the severity range, impaired nonword reading was the last symptom to resolve. This would also neatly explain the symptom complex that accompanies semantic errors; as the most severe symptom they should always be accompanied by all of the less severe symptoms. Thus, according to this classification, phonological dyslexics may actually be mild deep dyslexics (see also Patterson, 1982). Support for this view comes from Lambon Ralph and Graham (2000), who provided an overview of 85 deep/phonological cases and observed that, in the majority of cases, deep and phonological dyslexics had lesions in the same areas but that the lesions for deep dyslexics tended to be more extensive than those for phonological dyslexics. Berndt, Haendiges, Mitchum, and Wayland (1996) conducted a study of 11 patients selected on the basis of a single left hemisphere cerebrovascular accident (CVA) coupled with impairment in nonword reading. Of the 11, 3 were categorized as

making a significant number of semantic errors, and these 3 patients were also at the bottom end of performance on nonword reading. This fits with what one would predict from the continuum hypothesis.

More recently, Crisp and Lambon Ralph (2006) conducted a similar study of 12 CVA aphasic patients, all of whom exhibited a lexicality effect in reading. This study had the specific intention of exploring the continuum between phonological and deep dyslexia. Whilst there was a clear overlap in the patients' deep/phonological symptoms, there did not seem to be any evidence for Friedman's (1996) idea of a strict symptom progression. Most of the patients exhibited a significant imageability effect, previously thought to be rare in phonological dyslexia, and all exhibited a parallel impairment on delayed repetition, suggesting that their deficits stemmed from generalized phonological damage. Crisp and Lambon Ralph suggested that, rather than using Friedman's continuum, a better way of classifying patients would be in terms of the degree of impairment to their primary reading systems (phonology and semantics). This conclusion was supported by the results of tests on additional nonreading phonological tasks. All patients exhibited deficits on these tasks that paralleled their reading deficits (e.g., lexicality effects in repetition). What is more, the level of their impairment on phonological tasks such as blending and segmentation correlated positively with their nonword reading accuracy.

Attempts to produce computational models of acquired phonological dyslexia are relatively rare in comparison to the other central dyslexias. Patient case series such as Berndt et al. (1996) or Crisp and Lambon Ralph (2006) indicate that a successful model of phonological dyslexia must be able to produce lexicality effects that vary with severity of the phonological deficit apparent on nonreading tests: Mild damage should give relatively small lexicality effects arising from impaired nonword reading and near-perfect word reading; moderate damage should give large lexicality effects arising from near-floor nonword reading coupled with slightly impaired word reading; severe damage should again give relatively small

lexicality effects, this time arising from impaired word reading and abolished nonword reading.

Harm and Seidenberg (1999) successfully explored the phenomenon of developmental phonological dyslexia. They trained a single pathway O → P network in two stages. First they trained the phonological portion of the network so that it learnt 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. The key insight from this model was that symptoms of phonological dyslexia would arise from generalized damage located within the phonological system itself. Whilst they successfully modelled varying severities of developmental dyslexia, none of their simulations were intended to produce the very large lexicality effects found in cases of pure acquired phonological dyslexia. In fact, until very recently, there were no reported parallel distributed processing (PDP) models of acquired phonological dyslexia that simulate lexicality effects of the required magnitude. (Harm & Seidenberg, 2001, presented a model of acquired phonological dyslexia, but the focus of the paper was on orthographic influences on reaction times, RTs, and lexicality effects were not reported.)

Welbourne and Lambon Ralph (2007) demonstrated that it is possible to model all features of phonological dyslexic reading but only when the effect of plasticity-related recovery was included in the model (Welbourne & Lambon Ralph, 2005a). They used a replication of Plaut et al. (1996) and showed that damage to the phonological portion of the model followed by a period of plasticity-related recovery could simulate the full range of lexicality effects found in two case series studies (Berndt et al., 1996; Crisp & Lambon Ralph, 2006). This network also maintained the ability to simulate symptoms of surface dyslexia when the semantic input to phonology was removed, as seen in semantic dementia patients (Graham et al., 1994; Woollams et al., 2007). This was the first time

that this double dissociation has been explicitly simulated within the triangle framework. The insight from this and previous modelling work (Welbourne & Lambon Ralph, 2005a, 2005b) is that changes to the network, as the system partially recovers after damage, can be key to understanding the final patterns of deficits exhibited in the chronic stage. The importance of the period of partial recovery following damage, corresponding to the "spontaneous recovery" period observed in patients, is also key to the present simulation.

On this view, learning in the models can be seen as reaching an equilibrium between computational resources and the demands/characteristics of the learning environment. Initial development occurs as the learning environment pressures the network to improve its performance; eventually the network reaches an equilibrium where its performance is no longer improving—mature performance. However, brain damage destroys this equilibrium, and so the network is again pressured to find a new equilibrium—chronically impaired performance. The notion of recovery and reorganization is key to the account of phonological/deep dyslexia that we present here, and it directly leads to the possibility that DOL may change over the course of recovery, driven by the pressure to find a new equilibrium performance.

Key features of the model

Our aim in this paper is to produce a single unified "triangle-based" PDP model that can simulate all of the phenomena described above. In the process, we also hope to demonstrate the computational plausibility of the idea that damage followed by recovery can lead to shifts in the DOL, which themselves contribute to the patterns of chronic performance that make up the central dyslexias. To do this, we have incorporated a number of key features from previous versions of the triangle model together with some novel features. These are detailed in the following sections.

Plasticity-related changes in division of labour
Here and elsewhere, we have argued that plasticity-related recovery is a key component to simulating

dissociations such as those found in the central dyslexias (Welbourne & Lambon Ralph, 2005a, 2006, 2007). In our previous simulation studies, recovery reflected the reoptimization of the remaining processing units and connections within the damaged system itself. However, none of these simulations included a realistic implementation of semantics. In the full model, two drivers of recovery become possible: reoptimization of the remaining processing system within the damaged domain and also a reoptimization of the division of labour (DOL). As originally evaluated by Plaut et al. (1996, Simulation 4), the pronunciation of words reflects a mixture of direct activation and additional support from word meaning. In the undamaged system, efficient reading leads to an optimization of the DOL in favour of the direct computation $O \rightarrow P$ for the majority of items. As noted above, however, the same settings after damage may no longer be optimal, and, through plasticity-related recovery, the model will move towards a new DOL balance to maximize overall performance. Indeed, Crisp and Lambon Ralph (2006) speculated that phonological-deep dyslexia might reflect an increased semantic contribution, which might help to stabilize the patients' core phonological impairment. By extension, in the face of unreliable/impaired semantic input, the reading system might shift in the opposite direction, favouring the direct $O \rightarrow P$ computation more than normal, as shown by the prevalence of LARC errors in semantic dementia (SD). The following simulations explore this hypothesis. Specifically, in the case of phonological dyslexia we expected that the DOL would change to favour $O \rightarrow S \rightarrow P$ as recovery progressed, while in the case of surface dyslexia we predicted that the $O \rightarrow P$ pathway would become more dominant.

Model architecture

The architecture of the model is shown in Figure 1. It consists of three layers of interconnected units labelled orthographic, phonological, and semantic. There are 18 visible units (units with specific training targets) in each of the layers, and these are used to encode the representations appropriate to each layer. In addition, the phonological and semantic layers each contain 50 hidden units, which are free to develop their own representations under the pressure of training. Units in the orthographic layer are connected to units in the phonological and semantic layers with a 30% probability of connection. In addition, the phonological and semantic layers are bidirectionally connected to each other, again with a 30% probability of connection. Units within the semantic and phonological layers also have local connections among themselves such that any two units within one of these layers has an 80% chance of being connected.

All units have unambiguous, single affiliation

As can be seen in Figure 1, the most obvious novel architectural feature is the migration (compared to previous versions of the triangle model) of the hidden units to reside in the same region as either the visible semantic or phonological units. As such, there are now "phonological" hidden units and "semantic" hidden units rather then a single layer of hidden units that maps between phonology and semantics. This allows us to define DOL very simply in terms of the consequences of severing all $O \rightarrow P$ connections versus severing all $S \rightarrow P$ connections.

Connectivity density and type

The second novel architectural feature stems from knowledge about cortical connectivity. In the vast majority of previous PDP models, every unit in one layer connects with every unit in the next (Plaut, 1995a, 1995b, 1996; Plaut & Kello, 1998; Thomas et al., 2001; Welbourne & Lambon Ralph, 2005a; Zevin & Seidenberg, 2006; Zorzi, Houghton, & Butterworth, 1998, and many more). However, this 100% connectivity is unrealistic; if every neuron in the brain were connected with every other, then the size of the brain would need to increase to a sphere with a radius of 10 km (Nelson & Bower, 1990; Plaut, 2002). A more realistic view of cortical connectivity would be that nearby neurons are connected with a high probability via intracortical connections, whereas distant neurons are connected more sparsely via white matter intercortical connections.
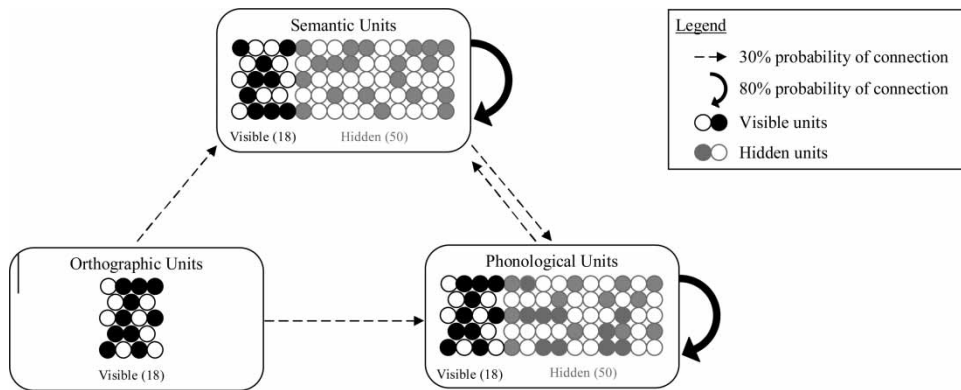
Figure 1. *Architecture of the reading model.*

Histological studies have demonstrated that this is certainly the case for mice and cats (Young, Scannell, & Burns, 1995), and it seems reasonable to assume the same for humans since high-density local connections are physiologically economic. Our model implements this neuroanatomical constraint by using two connection densities. Connections between units in the same system occur with a probability of .8 (each unit is connected, on average, to 80% of the other units in the system), whereas connections between units in one system and units in a different system occur with a probability of .3. It is this variation in connection density that encourages specialization of function. If all the units were fully interconnected then there would be no functional difference between the semantic or phonological hidden units. However, the strength of the intracortical connections relative to the intercortical ones means that it is more efficient for units in the phonological system to do phonological processing and units in the semantic system to do semantic processing (see Plaut, 2002, for the original development of this idea). As well as being relatively sparse, intercortical connections were restricted to only having excitatory connection weights. This facilitates an increased semantic contribution to reading as well as being neurophysiologically realistic: Intercortical connections (which only occur between pyramidal cells) are universally excitatory in nature (Braitenberg & Schüz, 1991).

Interleaved training

Another essential feature of the model's training regime was the use of interleaved training throughout the initial training and recovery periods. Various previous models of reading and other language activities have used this approach to some degree (Harm & Seidenberg, 1999, 2001, 2004; Joanisse & Seidenberg, 1999; Kello & Plaut, 2003). In doing so they permit (a) a simulation of the simultaneous development of various language activities, (b) a resultant model that can simulate all of these domains within the same framework (rather than having models for each task), and (c) the influence of continued exposure after damage on plasticity-related recovery.

Key features incorporated from previous models

We have already highlighted three important key features that have contributed to the success of previous models, which are adopted in the present simulations:

1. The development of a division of labour between semantics and phonology when reading aloud (Plaut et al., 1996); this is implemented in our model by using target semantic representations that are required to be activated alongside phonology when reading.
2. The insight that a lesion centred around a single location can produce both visual/phonological errors and semantic errors (Plaut & Shallice, 1993); this is implemented and

extended in our model; we show that lesions to the phonological system can produce symptoms along a continuum from phonological to deep dyslexia.

3. Plasticity-related recovery after damage (Welbourne & Lambon Ralph, 2005a, 2007).

Additionally, there are two features used in our models that are common to many other single word reading models but are nevertheless critical to our simulations:

1. Componential, vowel-centred representations for orthography and phonology (Plaut et al., 1996; Zorzi et al., 1998). These maximize the network's ability to generalize from the pronunciation of one word to another.
2. Quasi-regularity of the mappings between orthography and phonology, which contrasts with the arbitrary mappings to semantics (Lambon Ralph & Ehsan, 2006; Plaut et al., 1996).

SIMULATION 1: NORMAL SINGLE WORD PROCESSING

The goal in this simulation was to demonstrate that the model could learn to process all the words in its vocabulary in the reading, repetition, speech, and comprehension tasks and generalize that knowledge to read nonwords. The fully trained network from this simulation was then used as the starting point for additional investigations into surface and phonological/deep dyslexia.

Network dynamics

The network operated in continuous time with the input states of each unit changing according to an approximation of the following equation:

$$\frac{dx}{dt} = \sum_i s_i w_{ij} + b_j - x_j \qquad (1)$$

where x_j is the net input to the receiving unit j; s_i is the output of the sending unit i; w_{ij} is the weight of the connection between unit i and unit

j; b_j is the bias of unit j. For simulation purposes this equation is approximated by dividing time into discrete ticks. Here we used six arbitrary intervals of continuous time, each of which was discretized into five ticks. Under this approximation, Equation (1) can be modified to give a specific value for the change in net input between two ticks:

$$\Delta x_j = \tau \left(\sum_i s_i w_{ij} + b_j - x_j \right) \qquad (2)$$

The value of the averaging time constant τ was .2, which corresponds to the inverse of the number of ticks per interval.

As is usual in PDP networks, each unit's output (s_j) was calculated on the basis of its input (x_j) using a sigmoid activation function with the following equation:

$$s_j = \frac{1}{1 + e^{-x_j}} \qquad (3)$$

Error on the target units was calculated from the difference between the output and the target using cross entropy (Hinton, 1989; Kullback & Leibler, 1951). The network was trained using back-propagation through time with a learning rate of 0.05 and a 0.9 momentum term applied with the modification that the premomentum weight step vector was bounded so that its length could not exceed 1 (a variant often known as Doug's momentum). Initial weights were set to random values varying evenly between −0.1 and +0.1 for the intracortical connections (S → S and P → P), while for the intercortical connections (O → S, O → P, S → P, and P → S), initial weights were set to small random, positive values varying evenly between 0.001 and 0.1. A small degree of weight decay was applied to prevent individual weights from becoming too large and to promote generalization—weight decay is a process whereby all of the weights in the network are reduced by a fixed factor (in this case 10^{-7}) after every weight update. All networks in this study were generated and trained using the LENS neural network simulator (Rohde, 2000).

Training stimuli

A total of 216 items were used to train and test the model; these items consisted of linked orthographic, phonological, and semantic codes. The orthographic codes were constructed to mirror CVC (where C is a consonant, and V is a vowel) words with each word consisting of six possible onsets, six possible vowels, and six possible offsets. Eighteen binary digits (three groups of six) were used to represent these words. Within each group, one unit was used to represent each possible letter. Phonological codes were formed from the orthographic code in one of two possible ways. For consistent items, the phonological code was a direct copy of the orthographic one; these items constituted most of the set (176/216). However, 20 items were chosen at random to have inconsistent $O \rightarrow P$ mappings, and for these items the vowel portion of the phonological binary code was rotated one place to the right (i.e., if the "on" bit was in Position 1, it was moved to Position 2, etc.; if it was in the last position, it was moved into Position 1). The consistent spelling of the phonology of these inconsistent patterns provided the basis for developing a subset of pseudohomophonic nonwords (nonwords whose phonological decoding would be a real word). In English the equivalent would be a pseudohomophone like BLUD where the pseudohomophone has the same phonology as a real word with a more consistent relationship between orthography and phonology. These pseudohomophones were not used in the training corpus but were reserved as pseudohomophone test items. A further 10 patterns with consistent $O \rightarrow P$ mappings were removed from the training set and were reserved for use as nonword test items. This left a training set consisting of 166 items with a consistent $O \rightarrow P$ mapping and 20 items with an inconsistent mapping.

This arrangement of orthographic and phonological representations captures two important features of English orthography: (a) For the majority of words the relationship between orthography and phonology follows a consistent pattern; for a small subset, however, this consistency is violated in the vowel portion of the word,[3] rendering these items inconsistent in terms of their input–output mappings; (b) the representations used for each are componential; onsets, vowels, and offsets each have their own sets of representations, which are built up to form the whole word. This helps to ensure that the network can capture consistencies between the pronunciations of different words (Plaut et al., 1996).

Semantic representations were generated for each of the 186 words in the training set. These representations were generated by P creating 18 bit random binary vectors subject to the following constraints:

1. The total number of "on" bits was always equal to four.
2. Half of the representations were restricted to have their "on" bits in the first nine digits, while the other half were restricted to have their "on" bits in the last nine digits (this ensures that the representations have a basic categorical structure).
3. Semantic representations were required to be unique.

Once generated, the semantic representations were randomly paired with existing orthographic/phonological representations. Although these representations are obviously not intended to reflect the complexity of meaning seen in natural language, this scheme captures the property that is critical for the current purposes—namely the arbitrary mapping between meaning and surface form that is true of the majority of words in most languages including English.

Test stimuli

Exploring the phenomena associated with all of the central dyslexias requires test stimuli that

---

[3] Although the vast majority of inconsistencies occur in the vowel portion of the word, there are some words in English with inconsistencies in the onset or coda (e.g., gaol vs. game, or cough vs. though). Our model is not large enough to allow us to include these types of rare inconsistencies. However, it seems very unlikely that this omission would be critical to our general findings.

manipulate lexicality, consistency, frequency, and imageability. Nonlexical test sets consisted of the 10 items that had been reserved as nonwords and the 20 pseudohomophones that had been created from orthographically consistent versions of inconsistent words. To create the required lexical test sets, the 20 inconsistent words were selected together with an additional 20 consistent words randomly selected from the training corpus. These were then divided into four further groups (5 items per group) on the basis of frequency and imageability. The frequency manipulation was achieved by scaling the error and error derivatives as a proxy for frequency. High-frequency test items were scaled by a factor of 10, while low-frequency test items were scaled by a factor of 3. This is effectively the same as scaling the learning rate for high-frequency items and has the same effect as using the frequencies 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 compressing the required training (see Plaut et al., 1996, for a fuller discussion of this issue).

Imageability was implemented in these simulations as a combination of "semantic richness" (Jones, 1985) and age of acquisition (AoA). Imageability is highly correlated with AoA, and some authors have argued that effects attributable to imageability are in fact caused by AoA. This has generated considerable debate (Ellis & Monaghan, 2002; Monaghan & Ellis, 2002; Strain, Patterson, & Seidenberg, 2002). Independent effects of both variables have been seen in normal readers' performance for inconsistent words (Monaghan & Ellis, 2002; Shibahara et al., 2003), although deconfounding of the two has not been attempted in the acquired dyslexia literature. Resolving this debate is beyond the scope of the current study. However, our pilot simulations suggest that richness of semantic representation and AoA (concrete concepts are generally early acquired) both contribute in a similar way to behavioural performance (resulting in slowed reading of inconsistent, low-imageability/late-acquired items in normal reading and less accurate reading of the same items following phonological damage).

Accordingly we adopted a neutral stance on this debate and incorporated both manipulations into our simulation of imageability. This also reflects the real-world situation where the vast majority of high-imageability words are also early acquired. To simulate semantic richness, learning rates were scaled up for high-imageability test words with this scaling restricted to connections linked to the semantic units in the network. (Note this is similar to the manipulation used to implement frequency. In both cases the magnitude of the weight updates is scaled up for the high-frequency/semantically rich items, but the frequency manipulation scales up error attributable to phonological and semantic units equally, whereas the semantic richness manipulation only scales that portion of the error attributable to the semantic units alone.) To simulate AoA, the training corpus was divided into two, with one half (including all the high-imageability test items) being designated early AoA, while the other half was designated as late AoA. Early-AoA items were introduced into the training set substantially before the late-acquired items.

Training procedure

Six different training tasks were used in this simulation: phonological maintenance (P → P); semantic maintenance (S → S); speech (S → P); auditory comprehension (P → S); letter sounds (O → P, no S); and reading (O → P, plus S). For the phonological and semantic maintenance tasks, the training patterns consisted of valid phonological or semantic codes (words). These patterns were soft clamped onto the relevant layer with a clamping coefficient of 0.9. Soft clamping allows the input units to respond more gradually to external (clamped) inputs with the asymptotic output value (in the absence of any other input) being given by the following equation:

$$Output = InitialOutput + ClampStrength.$$
$$(ExternalInput - InitialOutput) \qquad (4)$$

The network was trained for 15 ticks, after which the inputs were switched off, and the network was required to maintain the correct activation for a further 4 ticks. (Targets were

introduced after the first 5 ticks.) For training of the spoken pathway, inputs were hard clamped to the input layer, and the network was trained for 20 ticks (again targets were only introduced after the first 5 ticks). Training in knowledge of letter–sound correspondences (analogous to phonics teaching in children) was administered by hard clamping one of the orthographic positions to a valid letter and training the network to activate just the correct phonology for the corresponding phonological position. For training in the reading task, the orthographic inputs were hard clamped, and the network was allowed to update freely for 15 ticks. After this, targets were applied to both semantic and phonological units, and the network was trained for a further 15 ticks.

Ten versions of the network were trained, each with different initial weights. Training was split into four phases, and within each phase there was the possibility of training on a number of different tasks (see Figure 2). In each phase, training proceeded in the same way. First one of the available tasks was probabilistically selected, then a batch of 40 items was randomly selected from the training corpus, and the network was trained on all items in the batch for the selected task (weight updates were applied at the end of each batch of training). In Phase 1, training was split evenly between learning the semantic and phonological maintenance tasks. In Phase 2, training on speech and verbal comprehension was added: Two thirds of the training episodes were devoted to learning the mappings for the spoken pathway and one third to maintaining knowledge of the valid semantic and phonological representations. In Phases 3 and 4, training on reading was introduced as well, one third of the training episodes were used to maintain knowledge of valid semantic and phonological representations, one third were used to maintain knowledge in the spoken pathway, one sixth were used to learn the letter–sound correspondences, and one sixth were used to learn reading. This phased structure of training was intended as a rough approximation to human development, where the very early stages of learning involve developing phonological and semantic representations followed by learning to use those representations to support speech and comprehension. Only once a reasonable proficiency in those skills had been achieved was training on reading introduced along with training on individual letter/sound correspondences (phonics teaching). AoA was implemented in the training by delaying the introduction of late-acquired items by one phase. Phase 1 training lasted for 150 epochs, Phase 2 for 40,000 epochs, Phase 3 for 50,000 epochs, and Phase 4 for up to 70,000 epochs. (These time periods were selected following pilot simulations, but the exact lengths of the training periods are not critical to the results.)

Phase 1 150 Epochs	Semantic Maintenance Task *(Early Items)*			Phonological Maintenance Task *(Early Items)*	

Phase 2 40,000 Epochs	Semantic Maintenance Task *(Early Items)*	Phonological Maintenance Task *(Early Items)*	Speech Task *(Early Items)*	Comprehension Task *(Early Items)*	

Phase 3 50,000 Epochs	Semantic Maintenance Task (All Items)	Phonological Maintenance Task (All Items)	Speech Task (All Items)	Comprehension Task (All Items)	Letter/Sound Correspondence Task	Reading Task *(Early Items)*

Phase 4 70,000 Epochs	Semantic Maintenance Task (All Items)	Phonological Maintenance Task (All Items)	Speech Task (All Items)	Comprehension Task (All Items)	Letter/Sound Correspondence Task	Reading Task (All Items)

Figure 2. *Showing how the model was trained in four different phases.*

Stopping criteria

It is usual when training neural networks to stop the training when generalization has reached asymptote. This can be achieved by periodic testing using an untrained set of items (Bishop, 1995, pp. 343–345). In reading, this would be equivalent to stopping when nonword reading accuracy was at maximum. However, our network is much more complicated than the kind usually considered in standard treatments, mainly by virtue of the fact that we are training on a number of different tasks simultaneously. In this situation, adopting a stopping criterion based on only one of the tasks could result in very poor performance on the rest. The issue is further complicated by the fact that for many of the tasks (speech, verbal comprehension, and written comprehension) there is no possibility of generalization as the mappings are arbitrary. To overcome this problem, we developed a stopping criterion based on a composite score of performance for the training tasks plus the nonword reading performance. This was calculated as the simple average of scores on the following tasks: single word reading; speech; verbal comprehension; written comprehension; nonword reading. The asymptotic value for this composite score was used to determine the end point for training.

Results

For all tasks, a correct response was defined as a response where the difference between the target and the actual output was smaller than 0.5 on all of the output units for which a target activation was provided. Table 1 shows the mean network performance on the various training and testing tasks at the end of the training period. The network was able to perform all of the training tasks to an accuracy of at least 96%. In reading it was very slightly better at producing the phonology from orthography than it was at activating

Table 1. *Average performance of the 10 networks in training and generalization tasks*

Task		Mean (%)
Phonological	Maintenance (repetition)	100
	Speech	97
	Single word reading	100
Semantic	Maintenance	100
	Auditory comprehension	96
	Written comprehension	96
Generalization	Nonword reading (excluding pseudohomophones)	92
	Pseudohomophone reading	97

the correct semantics (100% vs. 96%). In addition, it was able to read nonwords to an accuracy of 92%, which is very close to the average human performance of 94% (Glushko, 1979).[4] It was more accurate at reading pseudohomophones than it was at reading nonwords (97% vs. 92%), perhaps implying that the pseudohomophones were benefiting from additional semantic support (see, Harm & Seidenberg, 2004, Simulation 17). However, over the 10 trial networks this difference was only marginally significant, $p = .076$, $t(18) = 1.498$. It is interesting to note that the model was also slightly faster at reading these stimuli then nonwords, consistent with the pseudohomophone advantage found in normal participants when reading mixed lists of pseudohomophones and nonwords (Reynolds & Besner, 2005).

RTs for word reading were calculated as the number of ticks that the phonological units required to settle into the correct pattern (all units within 0.5 of their targets). Settling was defined as having the same binary activation pattern for 2 successive ticks after the end of the grace time. To ensure sufficient data for reliable results, the 10 networks' RT times were sampled every 1,000 epochs for the last 40,000 epochs of training (the networks were performing near asymptote throughout this period). Table 2

---

[4] This result depends upon a very stringent criterion whereby only one possible pronunciation of each nonword is allowed; if we had accepted variations where the vowel had been pronounced to match an inconsistent word than the accuracy rate would have been very close to 100%.

Table 2. *Mean reaction time performance for the 10 networks on words and nonwords*

Frequency	Consistency	Imageability	Mean	SD	MSE
High	Inconsistent	High	17.48	1.20	0.03
High	Inconsistent	Low	17.91	1.47	0.03
High	Consistent	High	17.00	0.05	0.00
High	Consistent	Low	17.01	0.31	0.01
Low	Inconsistent	High	17.84	1.48	0.03
Low	Inconsistent	Low	18.74	1.80	0.04
Low	Consistent	High	17.01	0.33	0.01
Low	Consistent	Low	17.08	0.75	0.02
Nonwords			17.65	2.35	0.05
Pseudohomophones			17.17	1.16	0.01

Note: Performance measured in ticks from presentation of input patterns.

shows the mean, standard deviation, and mean square error values for all of the stimulus groups.

The RTs for the different test sets were submitted to a $2 \times 2 \times 2$ analysis of variance (ANOVA), where the variables were frequency, consistency, and imageability. The RTs for nonwords were compared for those of low-frequency inconsistent words and high-frequency consistent words (both collapsed across imageability) using a t test.

The standard effects of frequency and consistency were all present. There was a significant interaction between frequency and consistency, $F(1, 16098) = 257$, $p < .001$, as well as significant main effects of frequency, $F(1, 16098) = 329$, $p < .001$, and consistency, $F(1, 16098) = 3,081$, $p < .001$. In addition, there was a significant three-way interaction between consistency, frequency, and imageability, $F(1, 16098) = 37$, $p < .001$. The form of this interaction can be seen in Figure 3 and parallels that observed by Strain, Patterson, and Seidenberg (1995), such that the imageability effect was most pronounced for the low-frequency inconsistent words. There was also a significant frequency–imageability interaction, $F(1, 16098) = 56$, $p < .001$, as well as a significant consistency–imageability interaction, $F(1, 16098) = 321$, $p < .001$, and a significant main effect of imageability, $F(1, 16098) = 404$, $p < .001$.

Mean nonword reaction times (17.6 ticks) significantly slower than those from high-frequency consistent words (17.0 ticks), $t(2, 682)$ = 14.0, $p < .001$. They were also slightly faster than reaction times to low-frequency inconsistent words (18.2 ticks), $t(4, 489) = -12.1$, $p < .001$, and slower than reaction times to pseudohomophones (17.2 ticks), $t(3, 171) = 10$, $p < .001$. This is similar to human performance where nonwords are read more slowly than consistent words, whilst pseudohomophones are read more quickly than nonwords (Reynolds & Besner, 2005). However, the phenomenon where the model reads regular nonwords significantly faster than low-frequency inconsistent words is not usually significant in human experiments where regular nonwords are read at a similar speed to inconsistent ones (Glushko, 1979).

Restricted intercortical connectivity and hidden unit "specialization"

One key architectural assumption of this version of the triangle model is that restricted intercortical connectivity will promote something analogous to cortical specialization. Put simply, we assume that the phonological hidden units will do phonological processing because it is more efficient for them to do so as they are more densely connected to the phonological output units than the semantic output units. To test this assumption, we trained another version of the model where there was no distinction between intercortical and intracortical connections (both were connected with a density of 60%). We then observed the difference in error
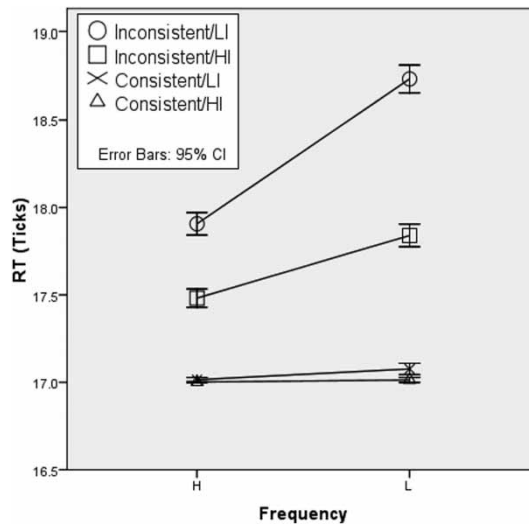
Figure 3. *Illustrating the three-way interaction between frequency, consistency, and imageability in the model's word reading reaction time (simulation of the "Strain" effect; Strain, Patterson, & Seidenberg, 1995). LI = low imageability. HI = high imageability. H = high. L = low. RT = reaction time.*

scores at the phonological and semantic output units arising from removing all of the connections from either phonological or semantic hidden units. Our prediction was that, where intercortical connections were restricted, damage to the connections from semantic hidden units would disproportionately affect semantic output, while damage to the connections from phonological units would disproportionately affect the phonological output. In the case where intercortical and intracortical connection densities were the same, we expected that there would be no significant difference in the error arising from damage to the different locations. Figure 4 shows the results of this comparison. As predicted, when the connection densities were homogeneous there was no significant difference in error arising from damage to connections from phonological or semantic units (all $ps >$ .05, all $ts < 1.5$). However, when the intercortical connections had a lower connection density than the intracortical ones, damage to connections from phonological hidden units resulted in a greater increase in error at the phonological output than did damage to connections from semantic hidden units, $t(19) = 4.35$, $p < .001$. Similarly, damage to connections from semantic

hidden units resulted in a higher error at the semantic output than did damage to connections from phonological hidden units, $t(19) = -3.17$, $p = .005$. In conclusion, these simulations demonstrate that relatively sparse intercortical connections encourage functional specialization in the connected units (see Keidel, Welbourne, & Lambon Ralph, 2010; Plaut, 2002, for a closely related idea).

Changes to division of labour during development

Harm and Seidenberg (2004) were interested in the DOL in terms of its contribution to the activation of semantics in normal reading. They used a lesion methodology where they assessed the accuracy of the semantic representation in the absence of either the O → S connections or the O → P connections. They showed that early in development the model is very dependent on O → P but later the O → S pathway also becomes important. In our case the focus for DOL investigation was on the activation of the phonological units rather than the semantic ones as we were concerned with reading aloud rather than reading for meaning.
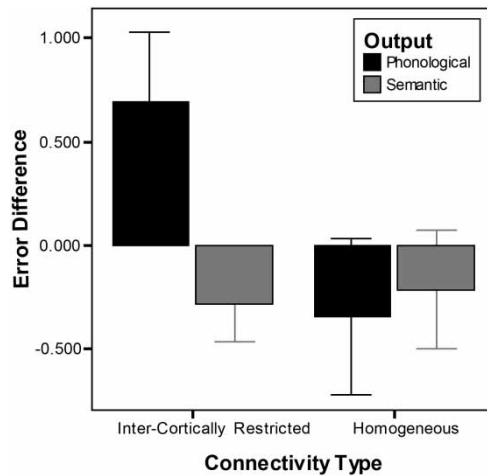
Figure 4. *The difference in error scores arising from selective phonological and semantic damage under two different connectivity regimes.*

Measuring DOL

We required a measure of DOL that would continue to be sensitive even when performance was very poor (after damage). Pilot investigations quickly demonstrated that using accuracy as a measure of DOL would not work for this situation. The method depends on removing the component of the system that is not of interest and measuring the accuracy of the remainder. We required DOL measurements in systems that had already been lesioned; further lesioning of these systems left accuracy levels at or near zero in most cases resulting in a major problem with floor effects. To counter this problem, we constructed a measure based on the reciprocal of the error score after lesioning. (Note that for these types of neural networks, error score is not % correct but a continuous measure of how close the network's output is to the ideal target response—high error scores indicate poor performance.) Using error scores to index DOL is advantageous because error scores remain sensitive even when accuracy is completely at floor. Using the reciprocal of the error score ensured that our measure was proportional to measures based on accuracy. Additionally, we wanted a measure that would allow us to compare the division of labour in networks with very different overall levels of performance.

To isolate the $O \rightarrow P$ contribution, we removed all the links between the semantic and phonological units and measured the error score in the phonological units (ErrorP). For the $O \rightarrow S \rightarrow P$ contribution, we removed the links between orthography and phonology and calculated the same error quantity (ErrorS). For each datum, the reciprocals of the two error scores were normalized to ensure that they summed to one. This was achieved by solving the following equation to find the normalizing constant.

$$\frac{Const}{ErrorS} + \frac{Const}{ErrorP} = 1 \quad Const = \frac{ErrorP + ErrorS}{ErrorP + ErrorS}$$

The resulting two measures allowed us to track changes in the relative contributions of the two pathways independent of the overall level of performance. The nature of these measures is that scores over 0.5 for a particular pathway indicate that that pathway is contributing more than 50% to the performance of the network.

In this simulation, we measured the network's DOL at 12 points every 10,000 epochs throughout development. Figures 5a and 5b show the results of this measurement in terms of the raw error scores and the normalized error reciprocals, respectively.

These results were consistent with what would be expected on the basis of Harm and Seidenberg
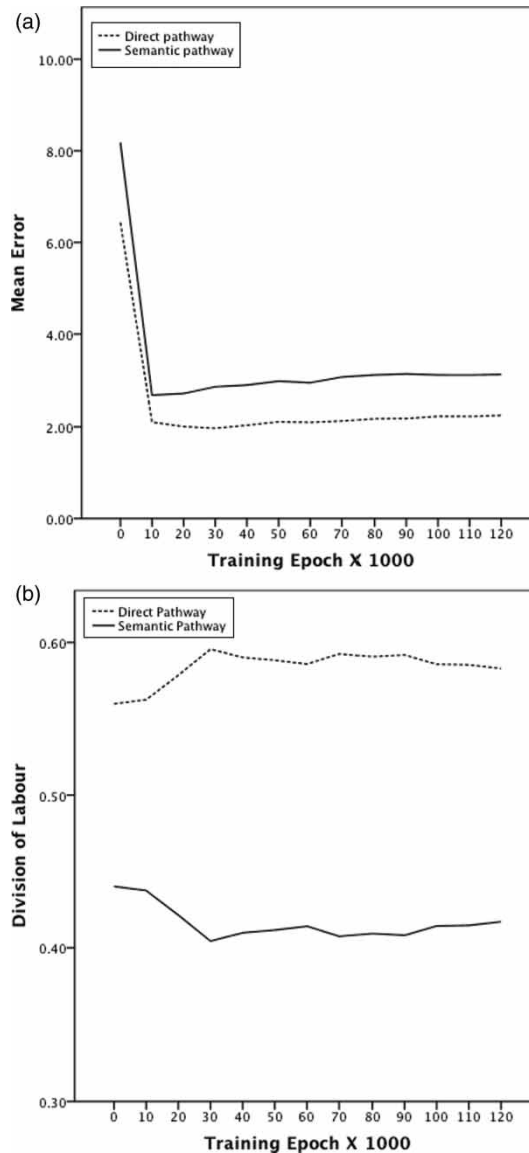
Figure 5. *(a) Error attributable to semantic pathway (O → S → P) and direct pathway (O → P) through development (O = orthography; S = semantics; P = phonology). (b) Division of labour between direct pathway (O → P) and semantic pathway (O → S → P) through development.*

(2004). Figure 5b shows that initially the division of labour moved strongly towards a position where the direct (O → P) pathway was doing the bulk of the work. This makes sense as this pathway requires mostly regular mappings, which are relatively easy to learn. However, after 30,000 epochs of training, the semantically mediated pathway (O → S → P) slowly began to increase its contribution to the correct activation of phonology, but even after 120,000 epochs of training its contribution was still much smaller than that from the direct (O → P) pathway. This is as

WELBOURNE ET AL.

expected given the quasi-regular nature of English orthography (Plaut et al., 1996).

SIMULATION 2: PHONOLOGICAL/ DEEP DYSLEXIA

The starting point for this simulation was the 10 fully trained reading models from Simulation 1. Unless otherwise stated, all methodological details for this simulation are the same as those for Simulation 1.

Patient data

We selected data taken from the case series conducted by Crisp and Lambon Ralph (2006) as target data for this simulation. These data included 12 participants with the symptoms of phonological or deep dyslexia. Participants were screened on the basis that they exhibited any one of the following symptoms in their reading aloud: (a) a lexicality effect, (b) an imageability effect, or (c) production of semantic paralexias. All the participants had acquired their dyslexia post cerebrovascular accident, but this diagnosis was not used as a method of selection. All were medically stable. There were 10 men and 2 women ranging in age from 40 to 83 years (mean age 59.4 years; $SD = 11.4$). Months post onset varied between 6 and 156 (mean 53 months; $SD = 47.2$).

Damage and recovery

Our hypothesis was that phonological/deep dyslexia occurs as a result of recovery after generalized phonological damage[5] (Crisp & Lambon Ralph, 2006: Note that all these reported cases of phonological dyslexia had at least three months post stroke recovery and usually much more than a year; Jefferies, Sage, & Lambon Ralph, 2007; Patterson & Lambon Ralph, 1999; Patterson & Marcel, 1992). In order to test this hypothesis in the model, we wanted to ensure that the location of

the damage was unambiguously "phonological" and that the damage also impaired the ability of the phonological system to recover. This is likely to be the case in the human brain where learning and representation are intrinsically linked. We achieved this by damaging a proportion of all the links (including internal links) that projected to or from any of the hidden or visible phonological units. The relearning capacity of the network was further impaired by adding an increased weight decay term to the remaining incoming connections within the damaged system. Four levels of damage were applied: Mild damage consisted of a 25% lesion to the phonological connections coupled with a weight decay of 0.000004 for the spared connections; moderate damage involved lesions to 35% of the phonological connections with a weight decay of 0.000008; severe damage was provided by a 45% lesion and weight decay of 0.000012; very severe damage was modelled by 55% lesions and weight decay of 0.000016. Each of the 10 trained networks was damaged separately 5 times at each damage severity, and the results were averaged. This averaging removes the danger that the results might by overly biased by a single unusual lesion, which is a potential problem for small networks (Bullinaria & Chater, 1995).

After the damage had been applied, the network was allowed to recover for 10,000 epochs, split into 10 periods of 1,000 epochs. After each of these 1,000 epochs, the performance of the network was tested. The balance of training tasks in the retraining period was adjusted slightly to reflect the likely experience of patents after damage, where most relearning experience is likely to focus on the basic language tasks of speech and comprehension. Accordingly 4/7 of the relearning episodes were devoted to relearning speech and comprehension, while 2/7 were devoted to maintenance of phonological and semantic representations. Only 1/7 of relearning episodes were devoted to relearning reading (compared to 1/3 in the original development).

---

[5] This may not be the only kind of damage that could theoretically lead to phonological/deep dyslexia, but it is by far the most common.

82

Results

Figure 6 shows the course of the network's recovery on the four basic language tasks of speech, spoken comprehension, reading aloud, and written comprehension after the various levels of damage. In all cases, the largest performance improvement occurs in the first half of the recovery period, although all cases continue to improve to some degree right up to the end of retraining. Performance on the reading aloud task is the best, reflecting the regular nature of the $O \rightarrow P$ mappings that dominate this task. It should be noted

that the performance on the comprehension tasks probably does not reflect the kind of scores one would expect to see on clinical "comprehension tests", as these assessments do not necessarily require a completely accurate activation of a semantic representation—the most common task used to assess comprehension is word-to-picture matching where a word must be matched to the correct picture selected from a group of distractors. Clearly, this task would be possible with an incomplete semantic representation provided that it was still sufficiently accurate to differentiate from the representations elicited by the distractors.
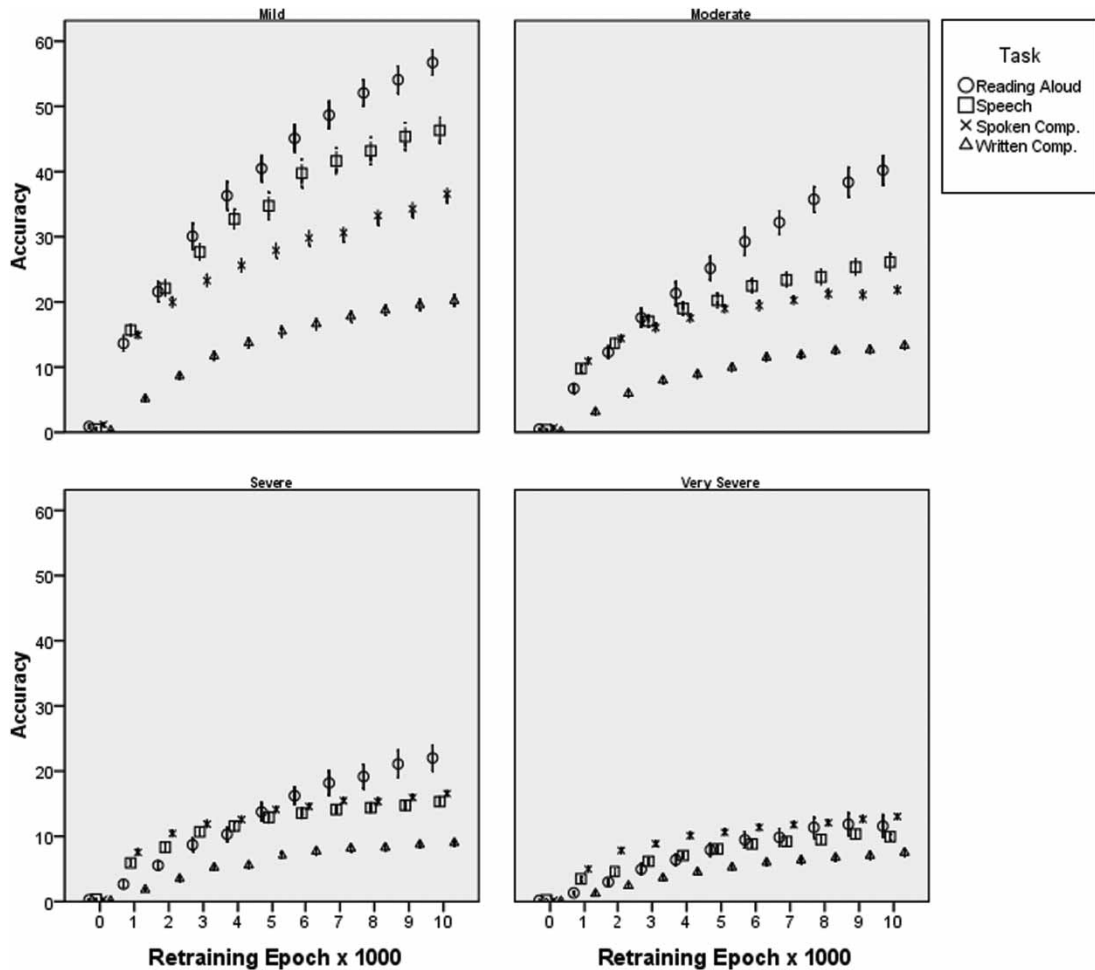

Figure 6. *Recovery of language function in the model following varying severities of damage. Comp. = comprehension.*

The most obvious effect of increasing damage on these tasks is that it reduces the ability of the network to recover. There is, however, another more subtle effect whereby recovery in reading aloud is more severely affected by increasing damage levels than are the other tasks. This can be seen in Figure 6 by comparing the mild and moderate damage cases, where recovery in reading considerably exceeds recovery in other language tasks, with the very severe case, where recovery in reading is similar to recovery in speech and spoken comprehension.

Overall the extent of the recovery is very considerable; after damage, the network performance recovers from 0 to between 10% and 50% of its undamaged function. This sort of spontaneous recovery trajectory is typical of what one observes clinically in patients presenting with aphasia following stroke (Pedersen, Jørgensen, Nakayama, Raaschou, & Olsen, 1995).

These data raise an interesting question concerning the appropriate point to start comparing the recovered model with patients. One might take the approach that the model should be allowed to reach asymptotic performance before a comparison is made, but this is problematic for two reasons: (a) The model never actually reaches a true asymptote; it just goes on improving by progressively smaller amounts; (b) in reality, patients are tested at many different points varying from as little as six months to many years post onset, and it is very unlikely that in all (or even most) cases their performance will have reached asymptote either. We adopted a pragmatic solution to this by electing to sample the network every 1,000 epochs after 5,000 epochs of recovery, continuing for 5,000 epochs. In this way, we created a sample where all networks have already recovered substantially but there is still some variation in the degree of additional recovery that could be expected. Our view is that by this method we achieve the most realistic match to the population of patients sampled in neuropsychological studies. However, we must acknowledge that this is an area of uncertainty—there is no foolproof way to ensure a completely accurate match on the basis of degree of recovery without detailed longitudinal patient data, which for obvious practical reasons is not currently (or ever likely to be) available.

Lexicality effects

The presence of a marked lexicality effect in reading accuracy is the cardinal symptom of phonological dyslexia. It is, therefore, essential that any model of phonological dyslexia be able to produce the same sort of range of lexicality effects as is found in patients. Figure 7 shows the range of lexicality effects found in the Crisp and Lambon Ralph (2006) case series compared with that found in the model. The patients' performance on word reading was as measured using the 40 high-frequency items from PALPA 31 (Psycholinguistic Assessments of Language Processing in Aphasia: Kay, Lesser, & Coltheart, 1992)—a set of polysyllabic words split by frequency and imageability. Nonword reading was assessed on all the items from PALPA 36 (Kay et al., 1992)—a set of nonwords varying in length between 3 and 6 letters. Word-reading performance in the model was measured by performance on the 10 items in the high-frequency consistent testing sets, while nonword performance was measured by performance on the group of 10 consistent nonwords. Network performance was sampled for the last 5,000 epochs of the recovery phase; accuracy rates were averaged across the 5 trials of each lesion severity for each of the 10 networks. This procedure resulted in 50 points for each lesion severity, with each point representing the average of 5 random lesions.

The picture from the model is very clear; there is a band of performance running from very poor word reading (10–30%) combined with almost completely abolished nonword reading, all the way up to near-perfect word reading with moderately impaired nonword reading (40–70%). There are not sufficient patients in the case series to map all regions of this band of predicted performance, but 10 out of the 12 patients clearly fall within the predicted band. The exceptions are A.B. and T.J. A.B. falls above the band because his nonword reading performance was too good relative to word reading, while T.J. falls below the band because his nonword reading performance
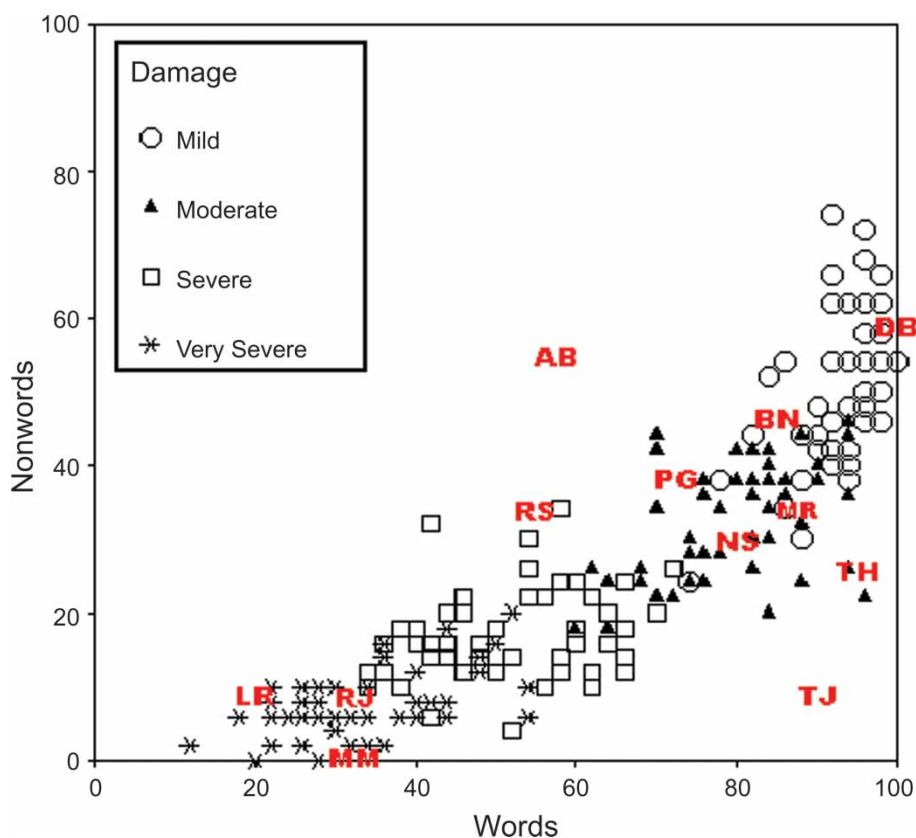
Figure 7. *Comparison of lexicality effects in the patients (indicated by initials) and the model. To view a colour version of this figure, please see the online issue of the Journal.*

was worse than predicted. A.B. could read 63.75% of nonwords and only 57.5% of words, so on these measures he had a slight reverse lexicality effect (on the measures used for screening he had a very small lexicality effect). The nearest point to A.B. from the modelling can read words with an accuracy of 70% and nonwords with an accuracy of 44%. T.J., on the other hand, could read words with an accuracy of 90% while only being able to read 8.33% of nonwords. The nearest point to T.J. from the models represented a word reading accuracy of 84% with a nonword reading accuracy of 20%. The fact that the patient sample has points on either side of the predicted performance band suggests that there is slightly

greater variability within the patient scores than is the case for the model. Despite this, the model has clearly passed the primary target for phonological dyslexia; it can reproduce a broad continuum of lexicality effects broadly matching patient data from mild to severe phonological dyslexia (Crisp & Lambon Ralph, 2006).

Pseudohomophone effects

In addition to effects of lexicality, phonological dyslexic patients often exhibit better reading of pseudohomophones than nonwords. The size of this pseudohomophone advantage can vary considerably (5–60%) depending on the language, the severity of the phonological impairment, and

choice of stimuli (Berndt et al., 1996; Patterson & Marcel, 1992; Patterson, Suzuki, & Wydell, 1996). In the Crisp and Lambon Ralph case series, the pseudohomophone effect was assessed on a matched set of 48 nonwords and pseudohomophones (Crisp, Howard, & Ralph, 2011). The average difference of pseudohomophones over nonwords was 13.9% ($MSE = 3.3\%$), which was significant for the group as a whole, $t(11) = 4.2$, $p = .001$. To gain an estimate of how the pseudohomophone effect varies with severity and thus match to the model, the patient data were split into two unequal groups, which corresponded to the moderate (9 patients: R.S., A.B., P.G., N.S., M.R., T.J., T.H., B.N., and D.B.) and very severe (3 patients: L.R., R.J., M.M.) categories in the model. Table 3 shows the average size of the pseudohomophone advantage for the four levels of lesion severity in the model together with data for the two levels of patient severity. The model produced significant pseudohomophone effects of between 10% and 20% depending on lesion severity. More severe lesions resulted in a smaller pseudohomophone effect. This matches the data from the Crisp and Lambon Ralph patient series in that not only is the average pseudohomophone effect very similar, but both patients and model show a similar reduction in effect size at the more severe level of damage.

Effects of imageability, frequency, and consistency

To test for effects of frequency, imageability, and consistency, accuracy rates for the eight groups of test words were submitted to a $2 \times 2 \times 2$ within-groups ANOVA. Data points were taken for each of the networks, sampled every 1,000 epochs over the last 5,000 retraining epochs and averaged across the five trial lesions. Table 4 shows the mean accuracy rates for each of the lesion severities, while Table 5 shows the results of the ANOVA. There are large and consistent main effects of all three variables: High-frequency items are read more accurately than low-frequency ones; consistent items are read more accurately than inconsistent ones; high-imageability items are read more accurately than low-imageability ones. There are also a number of significant two-way interactions. On closer inspection, however, these would seem to be artefacts of ceiling/floor effects. Evidence for this view comes from the fact that the direction of the interaction invariably changes as the severity of the damage increases. So for the frequency–consistency interaction arising with mild damage, there is a larger consistency difference for the lower frequency words (high frequency: consistent = 92.8, and inconsistent = 78.1, low frequency: consistent = 69.9, and inconsistent = 42.63); with moderate and severe damage, the interaction almost disappears, only to reappear, reversed, at very severe levels of damage, where there is a larger consistency difference for the higher frequency words (high frequency: consistent = 33.1, and inconsistent = 19.3, low frequency: consistent = 15.2, and inconsistent = 5.3). This pattern, where the greatest performance difference occurs for the pair of

Table 3. *Differences between nonword reading and pseudohomophone reading assessed at different levels of phonological damage severity in patients and model*

Lesion severity	NW	PH	PH advantage	MSE	t *value*	df	p
Mild	50.3	70.7	20.4	1.2	−13.6	49	<.001
Moderate	32.5	52.3	19.8	1.2	−16.3	49	<.001
Severe	17.2	30.9	13.7	0.8	−16.3	49	<.001
Very severe	7.6	16.8	9.2	0.7	−13.5	49	<.001
Patients (moderate)	28.5	44.4	16.0	4.2	3.8	8	.005
Patients (severe)	1.4	9.0	7.6	1.8	4.2	2	.053
Patients (all)	21.7	35.6	13.9	3.3	4.2	11	.001

Note: NW = nonword percentage correct; PH = pseudohomophone percentage correct; PH advantage = PH − NW; *MSE* = standard error of PH advantage.

Table 4. *Mean accuracy rates of the model for word sets varying on frequency, consistency, and imageability across varying levels of phonological damage*

Word set			Mild		Moderate		Severe		Very severe	
Freq.	Consist.	Image.	Mean	SD	Mean	SD	Mean	SD	Mean	SD
High	Cons.	High	94.8	5.2	83.9	10.2	60.6	13.0	41.6	12.8
High	Cons.	Low	90.8	6.4	75.9	10.4	43.1	11.1	24.6	9.6
High	Inc.	High	84.3	8.7	64.2	12.4	41.3	11.2	26.0	13.0
High	Inc.	Low	75.2	13.3	50.5	13.3	25.7	10.0	12.5	7.3
Low	Cons.	High	78.3	12.0	59.8	11.2	30.5	11.1	16.9	10.9
Low	Cons.	Low	64.7	10.3	41.5	11.5	22.6	7.9	13.4	6.7
Low	Inc.	High	55.0	15.2	30.2	14.8	12.1	7.5	7.1	5.9
Low	Inc.	Low	30.1	12.3	14.2	6.2	7.8	5.1	3.4	3.8

Note: Mean accuracy rates in percentages. Freq. = frequency. Consist. = consistency. Image. = imageability.

Table 5. *Results of Frequency × Consistency × Imageability ANOVA based on the networks' accuracy scores after varying severities of phonological damage followed by recovery*

	Mild		Moderate		Severe		Very severe	
Variable	$F_{(1, 49)}$	p	$F_{(1, 49)}$	p	$F_{(1, 49)}$	p	$F_{(1, 49)}$	p
Freq.	1,347	<.001	1,551	<.001	768	<.001	375	<.001
Consist.	427	<.001	519	<.001	250	<.001	149	<.001
Image.	236	<.001	190	<.001	218	<.001	158	<.001
F × C	66	<.001	8.8	.005	<1	ns	4.7	.034
F × I	63	<.001	10.3	.002	41	<.001	102	<.001
C × I	21	<.001	1.11	ns	3.9	.052	1.8	ns
F × C × I	3.2	.079	4.5	.04	<1	ns	1.4	ns

Note: ANOVA = analysis of variance. Freq./F = frequency. Consist./C = consistency. Image./I = imageability.

scores whose average is nearest 50%, is repeated for all of the two-way interactions. It is also likely to be responsible for the marginally significant three-way interactions found in the mild and moderate cases (all other three-way interactions were nonsignificant).

Figure 8 shows the magnitude of these main effects across the range of damage severities. At low levels of damage, the frequency effect is the largest with a difference of about 30% in performance on high- and low-frequency words, the consistency effect is somewhat smaller with a 21% performance difference, while the imageability effect is smallest with a 13% difference. As the damage severity increases, the effect sizes even out; for very severe damage the effects are all

roughly equal in size with performance differences of between 10% and 16%.

Unfortunately, it was not possible to make a perfect comparison between these results and those of the patients because there were no neuropsychological tests that manipulated frequency, imageability, and consistency simultaneously. Indeed, such sets are difficult to construct in English even for testing normal subjects (Ellis & Monaghan, 2002; Monaghan & Ellis, 2002; Strain et al., 1995, 2002). However, it was still possible to make a reasonable comparison by using data from a set of 96 words that varied imageability and frequency (Assessment 1a, from Crisp & Lambon Ralph, 2006) together with another set that varied consistency within low-
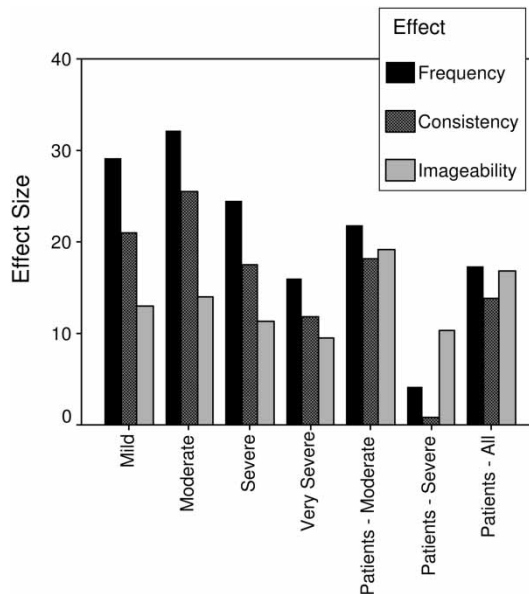
Figure 8. *Effects of frequency, consistency, and imageability in phonologically dyslexic patients and the model after phonological damage and recovery. The frequency effect is calculated by subtracting the accuracy rate for all low-frequency stimuli from the accuracy rate of all high-frequency stimuli. Consistency and imageability effects are calculated in a similar manner.*

frequency words (Assessment 3a also from Crisp & Lambon Ralph, 2006). These comparisons revealed that, overall, the patients showed a main effect of imageability, $F(1, 11) = 40.8$, $p < .001$, as well as a main effect of frequency, $F(1, 11) = 12.5$, $p = .005$: There was no significant interaction, $F(1, 11) = 2.5$, $p = .106$. They also showed a significant effect of consistency,[6] $t(11) = 4.41$, $p < .001$, one-tailed. The sizes of these effects (measured as the difference between adjacent levels of frequency and imageability) are also shown in Figure 8. For the 9 moderately impaired patients, the best match is with the severely damaged model. At this level of damage, all of the effects are of similar magnitude (approximately 20%) with the frequency effect slightly larger than the imageability and consistency effects. This is similar to the model's performance with severe

damage, although the imageability effect is somewhat smaller in the model than in the patients. For the 3 severely impaired patients the match is not quite as close. Like the model, the magnitude of all the effects has reduced, with the reduction more marked for frequency and consistency than it is for imageability. In the patients, however, the frequency and consistency effects have almost disappeared, whereas for the model they are still present, albeit at a lower level.

Correlations with nonreading tasks

A key finding from the Crisp and Lambon Ralph (2006) case series was that nonword reading accuracy correlated positively with performance on other phonological tasks such as blending and segmentation (Patterson & Marcel, 1992). In the model, the only purely phonological task is the

---

[6] This is not traditionally associated with phonological dyslexia. However, although it is not often reported, phonological dyslexics do often exhibit consistency effects. A reanalysis of data from Berndt et al. (1996) reveals that 9 out of 10 of the patients in the series showed more accurate reading of consistent than of inconsistent 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, $t(9) = 2.32$, $p = .023$, one-tailed.

phonological maintenance task—which is rather simple and does not allow us to model phonological awareness tasks directly. However, it was possible to examine correlations between error scores on the model's ability to maintain phonological activation and nonword reading performance. The prediction would be for a negative correlation; increased phonological error scores should go with decreased nonword reading accuracy. This proved to be the case ($r = -.86$, $p < .001$).

Analysis of error types

In addition to exploring the performance similarities between the model and the patients in terms of accuracy rates, it is important to establish whether the two produce a similar pattern of errors. The error types that could theoretically be produced by the network were analysed and then matched to the coding scheme used to analyse errors in the patients. Table 6 details the error categories produced by this process, together with their criteria. One difficult area for the model was the distinction between nonword and no-response (omission) errors. The nature of these models is such that some activation of the target units will always be produced. It would have been necessary, therefore, to come up with a criterion to decide whether a response was a nonword or an omission. For instance, one could choose a minimum threshold of activation below which the response would be regarded as an omission—if a high threshold away from 0.5 was chosen, then one would end up with lots of omission errors and few commissions. Alternatively, if set low it would generate many commissions and few omissions. Because this would be an essentially arbitrary decision, these errors were grouped together. Table 7 shows how the patients and the model compare on these criteria. Model errors are measured for the nine word and nonword test sets, while patient errors are measured on PALPA 31 (which varies imageability and frequency for the target words) and PALPA 36 nonwords. It is clear that, for both patients and model, the most likely error when reading a word is an omission or a nonword response, with most of the remaining errors being visual/phonological in nature. When reading nonwords, visual/phonological errors

become more frequent (for both patients and model).

As well as producing nonword/omission errors, the network also produces some LARC errors. The quantity of these errors is very sensitive to lesion severity; at mild levels of damage, 21.9% of all errors responses are of this type, but this rapidly reduces with increasing severity so that for severe levels of damage only 4.3% of errors are LARCs. These types of error are normally associated with surface dyslexics, where they are by far the most frequent form of error. Classically, one would not expect to see any of these errors in phonological dyslexia, but this may be partly due to lack of testing. In the Crisp and Lambon Ralph case series, only 4 of the patients were tested on sets of words designed to elicit LARC errors (Strain et al., 1995), and these patients did make LARC errors at a rate of 10.8%.

Visual/phonological errors are the model's second most frequent error type when reading words (between 7% and 15%, vs. 26% for the patients). For both patients and the model, the overall percentage of these errors declines with increasing severity. In addition, the model also makes some visual/semantic errors and a small number of pure semantic errors. The patients also make semantic errors, but for the severe group, they occur at much higher rates than in the model. (Note that each of the patients in the severe group made more semantic errors than any of the patients in the moderate group.) Before drawing any conclusions from these error rates, it is important to consider the possibility that the rates of lexical errors in the model may merely reflect the rates that would occur from a random sampling of the corpus (Ellis & Marshall, 1978; Plaut & Shallice, 1993). Over the whole corpus, the chance ratio of possible lexical errors was as follows: visual/phonological errors 6.4%; semantic errors 7.9%; visual/semantic errors 0.6%; other errors 85.1%. The actual distribution is very different: visual/phonological errors 73%; semantic errors 1.7%; visual/semantic errors 9.3%; other errors 15.9%. Comparing these expected chance ratios with the actual distributions from all runs of the model (using a chi-squared

Table 6. *Criteria for categorization of errors in reading*

Error type		Model criteria	Patient criteria
Word	Visual/phonological	Response is a word that contains at least 66% of the orthography/phonology of the target word and is not semantically related.	Response contains at least 50% of the orthography/phonology of the target word and is not semantically related.
	Semantic	Response is a word that differs by no more than two semantic features from the target, but contains only 33% or less of the target phonology.	Response is a semantically related word that is not also visually/ phonologically related.
	Visual/semantic	Response differs by no more than two semantic features and contains at least 66% of the phonology of the target word.	Response is semantically and visually/ phonologically related to target. This includes errors that would normally be categorized as derivational.
	Other lexical	Response is a word that is unrelated to the target.	Any other single word response.
	LARC	Response to an inconsistent word replaces the vowel with the more consistently used pronunciation.	LARC errors in response to inconsistent words[a]
	Omission/NW	Response is not a word.	Response is a nonword or no response.
Nonword	Visual/phonological	Response is a word that contains at least 66% of the phonology of the target.	Response contains at least 50% of the orthography/phonology of the target word and is not semantically related.
	Other lexical	Response is a word that is not phonologically related to the target.	Any other single word response.
	Omission/NW	Response is a nonword that was not the target nonword.	Response is a nonword that was not the target nonword or there is no response.

Note: NW = nonword. LARC = legitimate alternative reading of the components.
[a]Data only available for 3/12 patients on a different word set.

Table 7. *Errors made to words and nonwords by the model (after phonological damage and recovery) and phonologically dyslexic patients*

Error type		Percentage of errors from model				Percentage of errors from patients		
		Mild	Moderate	Severe	Very severe	Moderate	Severe	All
Word	Visual/phonological	14.6	13.6	10.9	6.8	32.7	19.5	25.5
	Semantic	0.3	0.4	0.5	0.3	0.7	9.8	5.6
	Visual/semantic	1.5	1.3	1.3	0.7	11.6	6.9	9.0
	Other lexical	2.5	3.4	3.4	2.9	0.7	2.9	1.9
	LARCs	21.9	15.8	8.7	4.3	n/a[a]	n/a[a]	n/a[a]
	Omission/NW	59.2	65.5	75.1	84.9	54.4	60.9	57.1
Nonword	Visual/phonological	37.6	30.2	17.2	11.7	59.9	54.4	58.0
	Other lexical	6.8	5.7	5.3	4.5	1.5	4.4	2.4
	Omission/NW	55.6	64.1	77.5	83.9	38.6	41.2	39.6

Note: NW = nonword. LARC = legitimate alternative reading of the components.
[a]The tests used in Crisp and Lambon Ralph (2006) were not designed to elicit LARC errors. However 4/12 of the patients, who had exhibited significant consistency effects, were tested on a set of words from Strain, Patterson, and Seidenberg (1995). Overall, these patients produced a LARC error rate of 10.8%. Error rates are expressed as a percentage of all incorrect responses.

test) reveals that there is a very significant deviation from chance for all severities: mild damage, $\chi^2 = 72,324$; moderate damage, $\chi^2 = 77,301$; severe damage, $\chi^2 = 61,124$; very severe damage, $\chi^2 = 38,126$, $p < .001$, for all severities.

The most obvious feature of the lexical error pattern is that the network is making many more visual/phonological errors than one would expect by chance alone, and, as a consequence of this, the rate of all other error types is reduced. To analyse the pattern of errors in more detail, we examined the ratio between each specific lexical error type and other nonspecific lexical errors.

The chance ratios can be calculated by dividing the chance rate of the specific lexical error by the chance rate of other lexical errors. These ratios were as follows: visual/phonological 0.0749:1; visual/semantic 0.00692:1; semantic 0.0928:1, respectively. The observed ratios for visual/phonological:other errors were 9.2:1, 5.5:1, 3.8:1, and 2.7:1 for the four levels of damage severity: between 123 and 36 times greater than would be expected by chance (all ps < .001). The observed ratios of visual/semantic:other errors were 1.2:1, 0.74:1, 0.49:1, and 0.31:1, respectively. These ratios were even more elevated over chance than the ratios for purely visual/phonological errors: between 173 and 44 times the chance rates (all ps < .001). The fact that the ratio for visual/semantic errors was higher than that for visual/phonological-only errors suggests that the network is also sensitive to semantic proximity. In the small number of cases where there is an available lexical error response that is both visually and semantically related to the target, then this word will be the most likely error for the network to make. The fact that visual/semantic errors are the preferred error type is interesting because this type of error is characteristic of deep dyslexia (Plaut & Shallice, 1993). Pure semantic errors are also characteristic of deep dyslexia, and these too occur at above chance rates in the model ($\chi^2 = 10.9$, $p = .001$).

Comparison with deep dyslexic symptoms
One of the key aims of this study was to attempt to elicit symptoms of all three central dyslexias.

Phonological dyslexia was thought to arise from generalized phonological damage, with more severe damage giving rise to symptoms characteristic of deep dyslexia. The cardinal symptom of deep dyslexia is the production of semantic errors. It is interesting that both the network and the patients have a tendency to produce some semantic errors. This supports the idea that there is a continuum between phonological and deep dyslexia (Crisp & Lambon Ralph, 2006; Friedman, 1996; Glosser & Friedman, 1990). For the model, this continuum appears both in the proportion of errors that are semantic in nature (which moves from .3 at mild damage to .5 at severe damage) and in the proportion of all responses that are semantic errors (0.2% at mild levels compared with 0.4% at severe levels). The pattern is the same for the patients; increased lesion severity results in a considerable increase in both proportion and total number of semantic errors. However, this increase is much greater than that observed in the model. For the 9 patients classed as moderate, the model gives a good match to the level of semantic errors, but the 3 severe patients made many more semantic errors than those produced by the model. However, absolute error rates are extremely dependent on the nature of the corpus, and it is not surprising that the model with a small corpus and a very restrictive definition of semantic errors makes fewer semantic errors than the patients. The key issue is whether the model makes more of these errors than could be expected given the nature of the corpus (Ellis & Marshall, 1978; Plaut & Shallice, 1993). This analysis has clearly shown that the model makes all the kind of errors found in deep dyslexic patients (visual, semantic, and visual/semantic) at levels well above those expected by chance alone.

In addition to the production of semantic errors, there are two other key deep dyslexic symptoms that any serious model of deep dyslexic reading must demonstrate. These relate to the property of imageability: Deep dyslexic patients are more likely to make errors when reading low-imageability words (Crisp & Lambon Ralph, 2006, showed that this was true throughout the deep–phonological dyslexia continuum); it is also
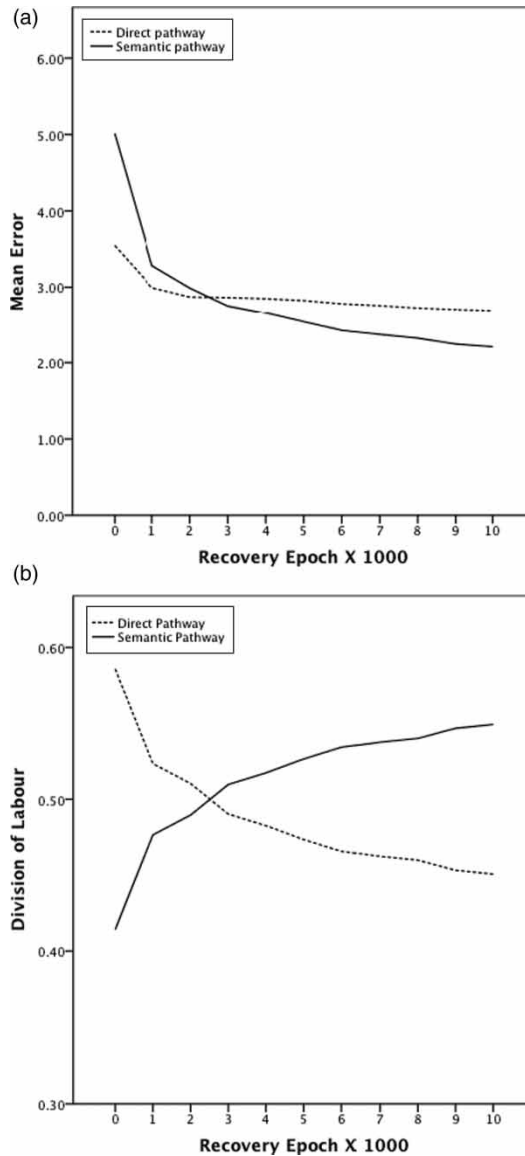
Figure 9. *(a) Error attributable to semantic pathway (O → S → P) and direct pathway (O → P) through recovery from phonological damage (O = orthography; S = semantics; P = phonology). (b) Division of labour between semantic pathway (O → S → P) and direct pathway (O → P) through recovery from phonological damage.*

likely that when they make visual/phonological errors the response will be more highly imageable than the target (Barry & Richardson, 1988; Nolan & Caramazza, 1982). We have already shown that the network was more likely to produce errors on low-imageability words than on high-imageability

ones. To test the influence of imageability on visual/phonological errors, we compared the imageability distribution of visual/phonological error responses for low- and high-imageability targets with the distribution that would arise from chance alone. For low-imageability targets there

92

was indeed a bias towards high-imageability responses, with a high-imageability response produced in 3,820 out of the 19,776 cases of errors on low-imageability words. This compares with an expected count of 2,126 (only 10.75% of the word corpus consisted of high-imageability words). High-imageability responses occurred 80% more often then one would expect by chance, $\chi^2(1) = 1,513$, $p < .001$. For high-imageability targets, this tendency was even more pronounced, with high-imageability responses produced in 231 out of 1,134 cases. This is 89% more often than would happen by chance, $\chi^2(1) = 109$, $p < .001$.

Analysis in terms of changes to division of labour

This simulation has demonstrated that damage to the phonological part of the network coupled with recovery produces behavioural effects in the model that closely resemble those observed in patients on the phonological/deep continuum (Crisp & Lambon Ralph, 2006). Phonological dyslexia can only occur in a system that is strongly biased toward semantic processing (resulting in lexicality and imageability effects: Crisp & Lambon Ralph, 2006). This requires a DOL that is also biased in favour of the semantic pathway. It may be, however, that this semantic bias is not a consequence of the damage itself but rather a result of reoptimization processes occurring in the recovery period. If so, we would expect that the DOL would move from direct pathway bias immediately after damage to a semantic pathway bias at the end of the recovery period.

To test this hypothesis, we selected the severely damaged networks and measured their division of labour at 10 points every 1,000 epochs throughout the recovery phase. As in Simulation 1, the O → P contribution was isolated by removing all the links between the semantic and phonological units, while the O → S → P contribution was identified by removing the links between orthography and phonology (these diagnostic lesions were performed after the initial phonological damage and recovery). Figures 9a and 9b show the results in terms of raw error scores and DOL

measurements, respectively. The results were exactly as anticipated. Immediately after damage, the system's DOL favoured the direct (O → P) pathway. In fact, the balance of DOL at this point was indistinguishable from the balance at the end of training. Phonological damage itself did not seem to have altered the DOL in any way. During the course of recovery, however, the DOL shifted decisively in favour of the semantic (O → S → P) pathway as the network reoptimized to make the best use of its remaining resources.

This is an extremely striking result; it is clear from Figure 9b that the recovery period allows for a plasticity-related functional reorganization that produces a cross-over in the DOL from greater dependence on the direct pathway to greater dependence on the semantically mediated pathway. It seems highly likely that this reorganization is also responsible for producing the lexicality effects that are typical of phonological dyslexia. However, it is also conceivable that the changes in DOL illustrated here are unrelated to the emergence of these effects. To eliminate this possibility, we compared the lexicality effect at the end of the period of recovery with the lexicality effects at the beginning (Figure 10).

This neatly illustrates the importance of the plasticity-related functional reorganization: Without it, there is no lexicality effect (the data fall along the diagonal representing equivalent word and nonword performance irrespective of the severity of damage); with recovery/reoptimization, there is a continuum of lexicality effects that corresponds well to the patients (the data are shifted considerably away from the diagonal to favour word > nonword performance).

SIMULATION 3: SURFACE DYSLEXIA ARISING FROM PROGRESSIVE SEMANTIC DAMAGE

Patient data

Surface dyslexia is characterized by a deficit in reading aloud low-frequency exception words,
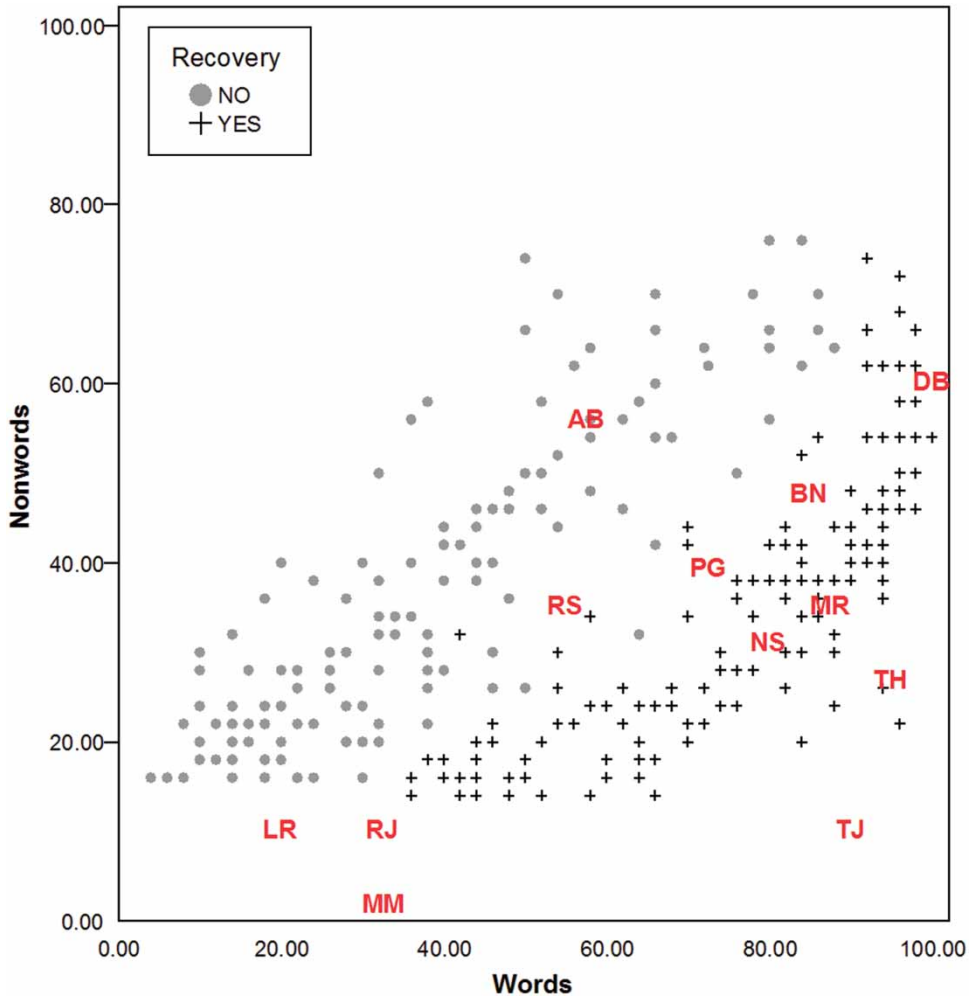
Figure 10. *The effect of phonological damage upon the magnitude of the lexicality effect with and without a recovery period allowing for plasticity-related functional reorganization. To view a colour version of this figure, please see the online issue of the Journal.*

combined with the presence of LARC errors. The vast majority of reported surface dyslexic cases come from a group of patients who suffer from semantic dementia,[7] a condition characterized by the progressive deterioration of semantic knowledge. We selected data of this type from the very large case series documented in Woollams et al. (2007). The bulk of the data set was derived from MemBrain, a patient database in Cambridge. The analysis included every observation of the "surface" reading list (Patterson & Hodges, 1992) recorded in MemBrain from a patient with an unambiguous clinical diagnosis of SD, provided that the reading data were accompanied by contemporaneous scores on picture naming and spoken word–picture

---

[7] We have elected to model surface dyslexia arising from progressive damage because that represents by far the largest number of reported cases. However, pilot simulations suggest that the model is also capable of reproducing the general pattern of surface dyslexic symptoms from acute damage followed by recovery.

matching (WPM). From this potential set, 1 patient was excluded because his naming scores were inflated by constant practice in naming these items as part of a rehabilitation study. Two other patients were excluded because their naming scores were at zero, and 1 patient's final score was removed owing to performance for the high-frequency regular words falling below 50%, suggesting a possible orthographic processing impairment. This selection procedure resulted in 88 observations from 43 patients. These were then supplemented with 12 observations from 8 SD patients (seen at a clinic in Bath) who were being tested on the same reading and semantic measures. The final data set consisted of 100 observations of from 51 SD patients, collected between 1991 and 2006.

Simulation targets

There were three key targets for this simulation: (a) The first was to demonstrate that progressive semantic damage could produce the symptoms of surface dyslexia. (b) The second was to explore the correlation between impairment in semantic tasks and reading accuracy. It is only with the advent of the very large case series study that the existence and quantification of these patterns could be demonstrated (Woollams et al., 2007). A key finding from this study was that the predictive value of a composite semantic score for reading accuracy is dependent on the consistency and frequency of the words in question. For low-frequency inconsistent words, 50% of the variance could be accounted for by semantic score alone, whereas for high-frequency consistent words, only 26% of the variance was accounted for. (c) The third was to test the hypothesis that the emergence of the surface dyslexic symptoms with progressive damage coincides with a gradual shift in DOL to favour processing through the direct ($O \rightarrow P$) pathway.

The composite semantic score used in the Woollams et al. (2007) case series consisted of the combined mean performance on naming and spoken word to picture matching. In these simulations, therefore, we used the task of mapping

from semantic to phonological representations as a proxy to naming, which was averaged with the verbal comprehension scores to give an equivalent of the patients' composite semantic score.

Semantic dementia-surface dyslexia: Combined damage and recovery

Unless otherwise stated, the details of this simulation are the same as those for Simulation 1. Starting from the 10 fully trained networks, semantic damage was applied in 30 separate steps (this gradual damage in the model mimics the progressive nature of semantic dementia). At each step, 1% of the remaining semantic links were removed while the remaining links had their weight decay term increased by 0.000001. After each episode of damage, the network was allowed to retrain for 50 epochs, after which performance on all tasks was tested. The balance of training tasks was the same as that for the final stage of initial training. This cycle of episodes of damage followed by retraining was repeated 20 times for each of the 10 networks. It was found that this procedure did not provide many examples of very mild damage (semantic scores greater than 75%)—this was rectified by running an additional simulation where the damage rate, weight decay, and learning rate were scaled down by a factor of 10. This additional simulation was sampled after 5 and 10 cycles of training.

Results

The above procedure resulted in a large set of points (8,400) in a space characterized by six dimensions corresponding to the scores on the five reading tests plus the composite semantic score. In order to compare the performance of the model with that of the patients, we matched each patient score with an average of 20 scores from the model selected on two criteria: (a) that the composite semantic score was within 1% of the patient's score, and (b) that the scores for low-frequency exception word reading were as close as possible to the patient's score. To facilitate the analysis of the effect of severity, we then
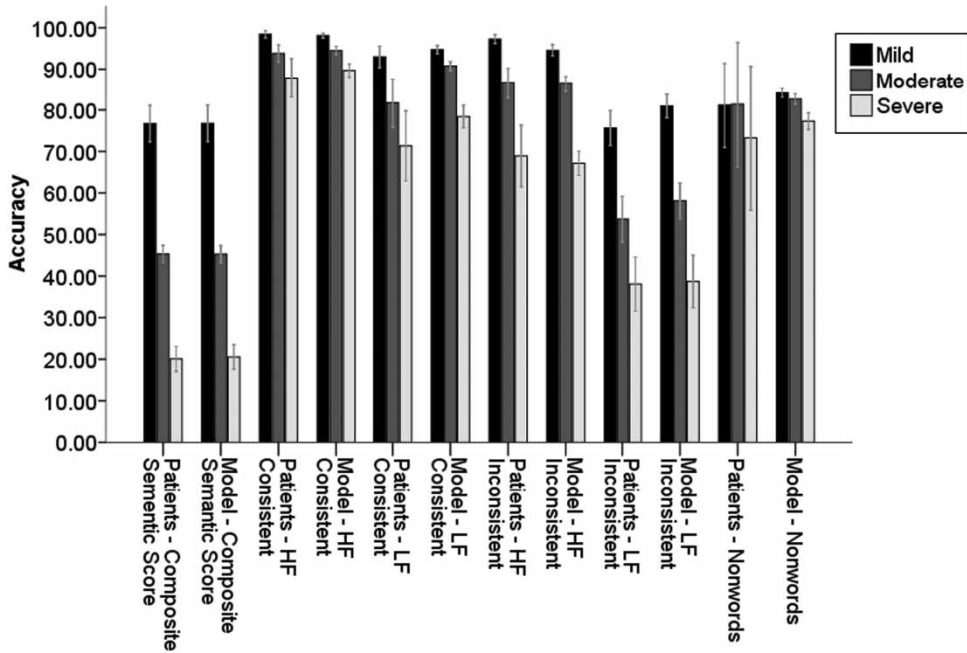
Figure 11. *Comparison of frequency/consistency effects in semantic dementia patients and the model. LF = low frequency. HF = high frequency.*

divided the scores into three equal groups (labelled mild, moderate, and severe) based on the patients' composite semantic scores.

Figure 11 shows how the reading task breaks down by word type and severity. Clearly there was a remarkably good match between the model and the patients right across the severity range.

One of the two key features of surface dyslexia is the presence of a frequency/consistency interaction such that low-frequency inconsistent words are disproportionately disadvantaged. To confirm that this key effect was present in our data, we submitted scores from each severity level to a 2 × 2 repeated measures ANOVA. In all cases, for both patients and model, there was a significant interaction as well as significant main

effects of frequency and consistency (all ps ≤ .001, one tailed). One interesting additional aspect of these data was that the end point for performance on low-frequency inconsistent words (for both the patients and the model) was still relatively high—approximately 40%. (Note, this was the average end point for both patients and model—there were individual examples in both data sets whose performance was lower than this, including examples with performance similar to patient K.T.[8]). This level of performance presumably reflects the ability of the phonological system when it is receiving no additional support from semantics. This contrasts with an end point for high-frequency consistent words of about 86%.

<hr>

[8] Based on the data reported by Plaut, McClelland, Seidenberg, and Patterson (1996) K.T.'s performance on reading was 26% for low-frequency inconsistent words and 100% for regular nonwords. A total of 85% of K.T.'s errors on reading exception words were regularizations (we would expect the LARC error rate to be similar). The closest matching point from the patient case series read low-frequency exception words with an accuracy of 29%, and 67% of errors were LARC errors. Unfortunately, no data were collected for nonword reading, but accuracy for regular words was 100%. The closest matching point from the model was 30% accuracy for low-frequency exceptions, 100% accuracy for nonwords, and a LARC error rate of 80%.

The second key feature of surface dyslexia is the prevalence of LARC errors when the target is an inconsistently spelt word. Figure 12 shows how the model and the patients compared in this regard. The model predicts that for all severities, LARC errors should make up between 56% and 62% of all errors; in general, the patients were making a slightly smaller proportion of LARC errors, usually between 40% and 54%, although for the mild group 74% of errors to low-frequency inconsistent words were LARCs.

To test the overall effectiveness of the model in predicting the patient data, we submitted all sets of actual and predicted standardized word reading scores to a single linear regression analysis. This confirmed that the model was a very significant predictor of the patient data ($R = .842$, $R^2 = .709$, $p < .001$).

One remarkable aspect of semantic dementia is the degree to which deficits co-occur. The reading scores from Woollams et al. (2007) correlated very strongly with a generalized semantic score constructed from two semantic assessments that involved no reading whatsoever (naming and spoken word to picture matching). Moreover, the slope of the resulting regression line varied with word type in a predictable manner. Low-frequency inconsistent words had the steepest slope; next came the high-frequency inconsistent words followed by the consistent low-frequency words and the consistent high-frequency words. To test whether this was also true of the model, the combined semantic score was used as a predictor for word reading accuracy for all of the different word types. Table 8 shows the results of this exercise together with the values obtained from the patients. The model showed a very similar pattern of results to that of the patients; the low-frequency inconsistent words had the highest B, followed by the high-frequency inconsistent words, then the low-frequency consistent words and the high-frequency consistent words. The nonword accuracy was not significantly predicted by semantic composite score. Importantly the confidence intervals for the value of B in the patients and the model were always overlapping.

Analysis in terms of changes to division of labour

To test the hypothesis that DOL would gradually move to become even more dependent on the direct (O → P) pathway, we tracked the DOL in the network as it was progressively damaged in semantics (again the DOL measurement procedure was the same as that used in Simulation 1 with the
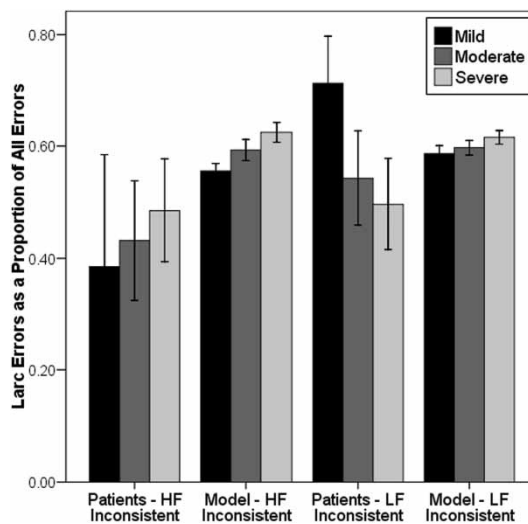

Figure 12. Comparison of proportion of LARC (legitimate alternative reading of the components) errors in the patients and in the model. LF = low frequency. HF = high frequency.

Table 8. *Results of regression studies using composite semantic score to predict single word reading accuracy*

	Patients		Model	
Stimuli	B	CI	B	CI
Inconsistent LF	.61	.49 to .74	.71	.61 to .81
Inconsistent HF	.45	.35 to .56	.461	.41 to .51
Consistent LF	.39	.26 to .52	.274	.233 to .315
Consistent HF	.19	.12 to .25	.146	.12 to .17
Nonwords	.16	−.18 to .49	.051	−.08 to .11

Note: Table shows unstandardized coefficients (*B*) and the confidence intervals (CIs) of those coefficients for patients and the model. LF = low frequency. HF = high frequency.

diagnostic lesions conducted on top of the lesions implementing progressive semantic damage). The expectation was that the semantic damage would tend to force the network to become more reliant on O → P processing, with the rebalancing of the DOL occurring in the retraining periods.

The results, which were as anticipated, are shown in Figures 13a and 13b for the raw error scores and DOL measure, respectively. These results conform to the pattern observed in previous PDP considerations of surface dyslexic reading in terms of a progressive reduction of "semantic" activation producing a relative increase on processing via the direct pathway (Plaut, 1997; Plaut et al., 1996; Woollams et al., 2007), and they agree with recent functional neuroimaging evidence showing that reading aloud in semantic dementia is more biased toward phonology than it is in matched control participants (Wilson et al., 2009).

Taken together with the results of the previous simulation, this makes a very strong case for the hypothesis that plasticity-related functional reorganization plays a substantial role in generating the patterns of behaviour seen in patients suffering from either surface or phonological/deep dyslexia. In the earlier case, phonological damage coupled with a period of recovery shifted the DOL to favour the semantically mediated pathway and produced symptoms of phonological and deep dyslexia. Here semantic damage interspersed with recovery shifts the DOL in the opposite direction and produces symptoms of surface dyslexia.

GENERAL DISCUSSION

This paper has described three simulations, conducted using a variant of the triangle model, where the specialization of phonological and semantic processing units emerged through training as a consequence of the model's internal connectivity structure. Simulation 1 demonstrated that this implementation of the model was able to reproduce a standard set of effects found in normal adult readers: good word reading; accurate nonword reading; an interaction between frequency and consistency in RTs; and semantic influences upon reading specifically for low-frequency inconsistent items (Strain et al., 1995).

Simulation 2 investigated the effect of phonological damage followed by recovery on the network's performance and probed its ability to reproduce data from the phonological–deep dyslexia continuum (Crisp & Lambon Ralph, 2006; Friedman, 1996). Detailed analysis of reading performance of the recovered networks revealed large lexicality and pseudohomophone effects coupled with additional significant effects of imageability, frequency, and consistency. As with patients, most errors were nonword/omission type errors. When lexical errors were made, they were most likely to be visual/phonological in nature. There was also an above-chance tendency for errors to be semantic or visual/semantic, particularly when the phonological damage was severe. This kind of pattern of performance is dependent on a division of labour where most of the work is being done via the O → S → P pathway. Analysis of the model as it recovered revealed that the required change in DOL is not a result of the damage itself, but emerges slowly entirely as a consequence of plastic changes driven by the relearning.

Simulation 3 used progressive damage to semantics to model surface dyslexia; the majority of cases of surface dyslexia occur in patients suffering from semantic dementia, which is a progressive degenerative disease that produces a specific and selective degradation of semantic knowledge (Hodges & Patterson, 2007; Patterson et al., 2006). Simulation 3 modelled this by interspersing repeated episodes of damage with exposure to the
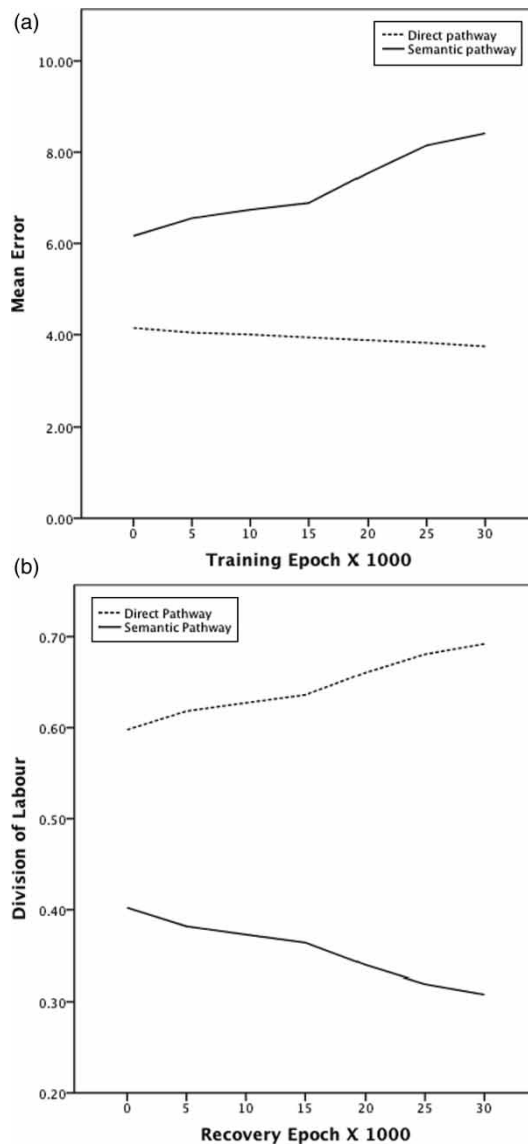
Figure 13. *(a) Error attributable to semantic pathway (O → S → P) and direct pathway (O → P) through progressive semantic damage (O = orthography; S = semantics; P = phonology). (b) Division of labour between semantic pathway (O → S → P) and direct pathway (O → P) through progressive semantic damage.*

original training environment. This resulted in a very good match to the data reported for the largest patient case series to date (over 100 observations of surface dyslexia in semantic dementia; Woollams et al., 2007): Reading accuracy on all items was reduced with increasing semantic damage, but the reduction was most pronounced

for low-frequency inconsistent words. Similarly, the percentage of responses that were LARC errors rose with increasing semantic damage. As reported in Woollams et al. (2007), the word reading accuracy rates were predicted by a composite semantic variable constructed from naming and verbal comprehension performance. For both

patients and the model, the slope of this graph was dependent on stimuli type, with the steepest slope found in low-frequency inconsistent words and the shallowest for high-frequency consistent words. Nonword accuracy was not significantly predicted by composite semantic score. Analysis of the DOL during this period of continuous deterioration shows that it gradually shifts to favour the direct (O → P) pathway driven by plastic changes occurring in the short recovery periods that intersperse each episode of damage.

The central hypothesis of this paper was that patterns of acquired dyslexia (the central dyslexias) depend on plasticity-driven functional reorganization that dynamically alters the division of labour (DOL) between the direct (O → P) pathway and the semantically mediated (O → S → P) pathway. This change in DOL was hypothesized to depend on a period of postmorbid recovery that corresponds to the spontaneous recovery stage seen in patients. These simulations have lent very strong support to this idea for four reasons: (a) The DOL in the model immediately after damage is almost indistinguishable from that in the undamaged model, suggesting that damage alone does not allow for a shift in DOL. (b) When the model is allowed to reorganize during a period of recovery, the DOL shifts in the hypothesized direction. (c) The patterns of performance only begin to resemble the central dyslexias when the DOL has had time to shift substantially from its position immediately after damage. (d) Under these conditions, the model is able to simulate symptoms of all three central dyslexias.

Key factors

The results of the simulations presented here pose an important question: What are the critical attributes of these simulations that are essential to modelling phonological/deep and surface dyslexia successfully? We suggest that there are three critical factors that were fundamental to the results of this simulation:

- Division of labour (DOL) between O→ P and O → S → P pathways.

- Plasticity-related recovery
- Triangle-based interactive architecture.

Below we consider the precise contribution of each of these factors in more detail.

Division of labour

A key consideration, when thinking about the system's performance, is division of labour. In normal reading (of English orthography), the natural division is for the direct pathway to do the bulk of the work; the regular nature of the O → P mappings (for English orthography) makes this the most efficient strategy (Plaut et al., 1996). To achieve surface dyslexic performance, the model shifts this balance even further towards the direct pathway to the point where there is no longer enough input from the semantic system to prevent LARC errors. In the case of phonological/deep dyslexia, the balance moves the other way so that the phonological output is more constrained by semantics than knowledge of the mappings between orthography and phonology. With mild degrees of phonological damage, this results in phonological errors on those items that are least supported by the dominant semantic system (nonwords and low-frequency, low-imageability words), while with more severe damage, phonological errors become widespread, and the dominance of the semantic system means that there is a tendency for some of these errors to be captured by words that are semantically related to the target—particularly if there is also some visual/phonological overlap. These simulations demonstrate that it is the changing division of labour that produces phonological–deep versus surface dyslexia. Critically, damage on its own was not found to produce noticeable shifts in the DOL; it is only when the network is allowed to reorganize itself after damage that the required shifts in DOL are generated.

These results cast considerable light on the processes underlying recovery and the production of performance dissociations in the model. The overt performance of this model, both pre and post damage, was driven by changes in the

DOL; it is only when the DOL had time to read-just that the model started to exhibit a corresponding behavioural dissociation. This would suggest that behavioural dissociations in neuropsychology do not always reflect subtraction alone (where the impaired performance implies damage to a critical subsystem), but can also arise from changes to the remaining systems post damage (countering the assumptions of "transparency" and "subtraction" used in some forms of traditional cognitive neuropsychology; Shallice, 1988).

Plasticity-related recovery

The original motivation for including plasticity-related recovery in the models stemmed from the observation that acute patients' performance changes rapidly in the six-month period immediately following damage and continues to do so over the subsequent months but at a diminishing rate (Lendrem & Lincoln, 1985). The period of recovery allows the damaged network to mimic this process, reorganizing its internal structure to find a new equilibrium where it can make the best use of its diminished resources. This has been explored to some degree in our previous work (Welbourne & Lambon Ralph, 2005a, 2007); however, those simulations did not include a semantic system and so were limited to showing how plasticity-related recovery might allow for reoptimization within the damaged system itself. The current simulations suggest that this recovery period does not just involve the damaged systems reoptimizing their performance, but also allows a global shift in the division of labour between the damaged and undamaged parts of the network.

Triangle-based interactive architecture

The triangle-based interactive architecture is helpful in that it allows reading to capitalize on connections already trained for other linguistic tasks that require mapping between phonology and semantics (see also Harm & Seidenberg, 2004). This allows the model to maintain a high accuracy in both sets of tasks whilst still allowing

good generalization to nonword reading. This occurs as a result of the pretrained connections between phonology and semantics, which strongly bias the model to learn to read by the direct pathway. This early bias towards the direct pathway is efficient because of the regular mappings. This early emphasis on learning the regular direct mappings is also beneficial for generalization performance. In addition, this type of architecture and training regime allows the model to explore associations between reading tasks and other linguistic tasks.

Additional model-specific factors

In addition to the three key factors listed above, there are two more model-specific factors, which, while they may not be essential to capturing all the target phenomena, have certainly been important in the development of this version of the triangle model and reflect known facts about neuroanatomy:

- Limited connectivity between different brain systems.
- Contrast between excitatory intercortical and more flexible intracortical connections.

These are discussed in more detail in the following sections.

Limited connectivity

This aspect of the model is necessary to allow differentiation between the different parts of the system. If the intercortical connections occurred with the same density as the intracortical ones, then there would be nothing to distinguish units in the semantic part of the system from units in the phonological part. The development of functional specialization is entirely driven by the fact that semantic units are more densely connected to other semantic units than to phonological units and vice versa. This structure then reflects the likely organization of connectivity in the brain, and a similar technique has been used to introduce functional specialization within the semantic system, allowing simulation of the optic

aphasias (Plaut, 2002), and also to explain the apparent paradox of emergent modularity and equipotentiality (Keidel et al., 2010).

Intracortical versus intercortical connections

As well as variations in connection density, another key distinction is that intercortical connections can only be excitatory. This reflects neurophysiology and has important functional consequences. For the current model, it ensures that the semantic system is adequately engaged in the process of reading aloud. This had a number of benefits in the model, including unmasking imageability effects and greater accuracy for nonword reading

Although it has clearly been successful in modelling surface and phonological dyslexia (PD), the current model was not able to match all of the features of deep dyslexia completely. Deep dyslexia and phonological dyslexia are very closely related. Recent studies have supported the notion of a phonologically based continuum between the two (with deep dyslexia as the end point: Crisp & Lambon Ralph, 2006; Friedman, 1996). Phonologically dyslexic patients demonstrated poor nonword reading, lexicality, pseudohomophone, and imageability effects. Deep dyslexic patients show the same features (in the context of worse reading overall); in addition they produce semantic paralexias. The model captures all of these characteristics, but the rate of semantic paralexias was lower than that for the patients. Despite this, the model was able to simulate the clear deep–phonological continuum, including the varying size of imageability, lexicality, and pseudohomophone effects. As per recent case series studies (Crisp & Lambon Ralph, 2006), the model demonstrated this behavioural continuum on the basis of increasing phonological damage. The model did generate more semantic errors than could be accounted for by chance alone. However, the absolute level of these errors was still quite low—less than would be seen in some more striking, "pure" cases of deep dyslexia. Of course, the absolute level of these errors is heavily dependent on the nature of the training

corpus and the definition of "semantic error", which between them determine the "chance" rate of such errors.

To help evaluate the quality of the model's match to deep dyslexic performance, it is instructive to compare our results with Plaut and Shallice's (1993) seminal deep dyslexia model. Their model used a smaller corpus but with a larger semantic vector, which allowed a more realistic category structure. Despite this, the absolute rates of semantic errors observed in their simulations were still relatively low—in the range of 0.1 to 1.2% (these rates are higher than expected by chance). Plaut and Shallice did not regard this quantitative difference as strong evidence against the adequacy of their model. Instead they argued that error rates can be greatly affected by a number of computational factors that are not central to the theory. Adopting this approach, we can make a strong case that the model has captured the essential quality of deep dyslexic performance (above-chance production of semantic and visual semantic errors). This view becomes stronger when one considers the additional features of deep dyslexic performance that the model produces: a large imageability effect coupled with a very strong tendency for lexical errors to be more highly imageable than the target word.

Other cases

This version of the triangle model has shown that it can simulate a wide range of patients suffering from any one of the central dyslexias. However, we must acknowledge that, while our patient data sets were large and representative of typical patterns of performance, there are some reported cases that might pose more of a challenge for the model to simulate. The challenge from these cases does not arise from their performance in the domain of reading (the model can cope with wide ranges of performance in this domain, including patients like K.T. who is widely regarded as representing the most extreme form of surface dyslexia), rather it comes from the very small set of patients whose reading disorder appears to be isolated from other linguistic impairments. The

triangle model does not specify a unique set of modules to support reading, rather it assumes that reading will be supported by the same set of brain systems that support visual and linguistic processes. The implication of this is that we should expect to see associations between reading deficits and deficits in other linguistic tasks that depend on the same primary systems. It is only with the advent of recent large-scale case series studies that the extent of these associations has started to be explored. However, it is now clear that there are very high correlations between semantic deficits and surface dyslexic symptoms (Woollams et al., 2007) and between phonological deficits and phonological dyslexic symptoms (Crisp & Lambon Ralph, 2006; Rapcsak et al., 2009), and these significant new data are captured by the current model. However, this means that there is a potential problem with the small number of cases that appear to contradict this view. These "counter" cases, or classical single dissociations, are examples of patients that display the "core" central dyslexia without the associated deficit in semantic or phonological processing or vice versa. With reference to the prevalence of surface dyslexia in semantic dementia, intact exception word reading has been reported in patient D.R.N. (Cipolotti & Warrington, 1995), E.M. (Blazely, Coltheart, & Casey, 2005), and initially in W.L.P. (Schwartz, Marin, & Saffran, 1979) and three cases (M.A., E.B., and M.G.) reported by Woollams et al. (2007). The opposite dissociation, of impaired exception word reading with intact semantic processing, has been reported in a single head injury case N.W. (Weekes & Coltheart, 1996) as well as two cases from Woollams et al. (W.M. and J.P.) Turning to instances of phonological dyslexia in the face of apparently intact phonological processing, these include: W.B. (Funnell, 1983), L.B. (Derouesné & Beauvois, 1985), R.R. (Bisiacchi, Cipolotti, & Denes, 1989), and, more recently, R.G. (Caccappolo-van Vliet, Miozzo, & Stern, 2004b), I.B., M.O. (Caccappolo-van Vliet, Miozzo, & Stern, 2004a), and J.H. (Tree & Kay, 2006). Our view is that at least some of these cases probably reflect premorbid individual differences: A tacit assumption of neuropsychology is that all cases were "identical" prior to their brain damage. However, this seems unlikely, and a more reasonable assumption would be that there are some premorbid differences in the division of labour in reading and that these differences may carry through to result in some measurable differences in the postmorbid performance. This idea can be used to account for the exceptional cases in the realm of surface dyslexia (Dilkina, McClelland, & Plaut, 2008; Plaut, 1997; Woollams et al., 2007)—readers who rely mostly on phonological processes will be likely to show less impairment in reading with mild semantic damage (e.g., cases M.G., E.B., and M.A. in Woollams et al. (2007)), while readers who rely heavily on semantics may show reading impairments while semantic deficits are still relatively mild (e.g., cases W.M. and J.P. in Woollams et al., 2007). It could also be used to account for cases of phonological dyslexia without additional phonological impairment—readers who premorbidly have a strong semantic basis may have quite poor nonword reading premorbidly so that even very mild phonological damage can produce phonological dyslexia. Simulating the effect of this kind of premorbid variability will be an important future challenge for this model. It should also be noted here that lexicality effects are not just found in patients with phonological impairment but also arise in those with primary visual/orthographic impairment (Rapcsak et al., 2009). The general explanation for this pattern mirrors that for phonological impairments—such that semantic feedback, or even plasticity-related semantic influences, boost visual processing of known words but offer little boost for nonword recognition.

Comparison with other approaches

This study has provided the first computational implementation of a model that captures aspects of all three central dyslexias. This is achieved as a consequence of plasticity-led changes to the DOL driven by postdamage recovery. At this point, it is probably worth pausing to consider

how this model relates to other approaches to modelling acquired reading disorders, principally the dual-route model (Coltheart, Rastle, Perry, & Langdon, 2001), the dual-process model (Zorzi et al., 1998), and the CDP+ model (Perry, Ziegler, & Zorzi, 2007). All of these models can produce a variety of features found in normal reading as well as some features of phonological and surface dyslexia. Both the dual-route cascaded (DRC) model (Coltheart et al., 2001) and the connectionist dual process (CDP+) model (Perry et al., 2007) can simulate phonological and surface dyslexia, while the dual-process model simulates surface dyslexia. Despite these similarities, there are two major areas where the current model differs from these other approaches: (a) the way in which they simulate damage; and (b) the ability to capture the continuum of performance between phonological and deep dyslexia.

The way in which they simulate damage

All three of the other models simulate damage by changing a parameter within the model in such a way as to reduce the influence of the damaged part of the model upon the output. Interestingly, this kind of damage is in effect a direct manipulation of the DOL between the lexical and sublexical routes in the models. If one puts aside for a moment the difference in labelling (lexical vs. semantic, sublexical vs. direct), then these models' simulation of acquired dyslexia are achieved in a similar way to the current model—that is, by changing the DOL to favour the undamaged system. The difference is that the current model implements a mechanism by which this can happen in a neural system—namely, destruction of some of the neural substrate followed by a period of plastic reorganization. This means that the model can explore how damage to a particular processing system (phonology or semantics) can naturally give rise to a change in the DOL through a plasticity-driven process of functional reorganization. This raises the intriguing possibility that models based on these principles might be able to make predictions about the likely course and extent of spontaneous recovery in aphasic patients or even provide guidance to therapists enabling them to choose a

therapy programme most suited to maximizing the functional reorganization.

The ability to capture the continuum of performance between phonological and deep dyslexia

The idea that there is such a continuum has gained considerable empirical support (Crisp & Lambon Ralph, 2006; Friedman, 1996; Glosser & Friedman, 1990), but this is the first time that it has been shown theoretically that the continuum can be simulated from increasing severities of phonological damage.

CONCLUSION

This study set out to provide a computational implementation of the hypothesis that plasticity-driven postmorbid relearning processes may dynamically alter the DOL during the period of recovery in such a way as to produce patterns of chronic performance that match patients with acquired reading disorders. It has demonstrated that the hypothesis provides an account of key aspects of phonological/deep and surface dyslexia. It has also provided evidence in support of the view that deep dyslexia forms a continuum with phonological dyslexia. The current model provides the only simulation of all three central acquired dyslexias that has appeared to date, capturing the critical patterns of performance across the full range of severity. Future challenges for the approach include implementation across a larger corpus of items and accounting for premorbid variability in DOL corresponding to different reading skills and styles.

REFERENCES

Andrews, S. (1982). Phonological recoding—Is the regularity effect consistent. *Memory & Cognition*, *10*(6), 565–575.

Barry, C., & Richardson, J. T. E. (1988). Accounts of oral reading in deep dyslexia. In H. A. Whitaker (Ed.), *Phonological processing and brain mechanisms*. New York, NY: Springer-Verlag.

Beauvois, M. F., & Dérouesné, 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.

Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford, UK: Clarendon Press.

Bisiacchi, P. S., Cipolotti, L., & Denes, G. (1989). Impairment in processing meaningless verbal material in several modalities: The relationship between short-term memory and phonological skills. *The Quarterly Journal of Experimental Psychology: A, Human Experimental Psychology*, 41(2), 293–319.

Blazely, A. M., Coltheart, M., & Casey, B. J. (2005). Semantic impairment with and without surface dyslexia: Implications for models of reading. *Cognitive Neuropsychology*, 22(6), 695–717.

Braitenberg, V., & Schüz, A. (1991). *Anatomy of the cortex* (Vol. 18). Berlin, Germany: Springer-Verlag.

Bub, D., Cancelliere, A., & Kertesz, A. (1985). Whole-word and analytic translation of spelling to sound in a non-semantic reader. In K. E. Patterson, J. C. Marshall, & M. Coltheart (Eds.), *Surface dyslexia* (pp. 15–34). London, UK: Lawrence Erlbaum Associates.

Bullinaria, J. A., & Chater, N. (1995). Connectionist modeling—Implications for cognitive neuropsychology. *Language and Cognitive Processes*, 10(3–4), 227–264.

Caccappolo-van Vliet, E., Miozzo, M., & Stern, Y. (2004a). Phonological dyslexia—A test case for reading models. *Psychological Science*, 15(9), 583–590.

Caccappolo-van Vliet, E., Miozzo, M., & Stern, Y. (2004b). Phonological dyslexia without phonological impairment? *Cognitive Neuropsychology*, 21(8), 820–839.

Cipolotti, L., & Warrington, E. K. (1995). Semantic memory and reading abilities: A case report. *Journal of the International Neuropsychological Society*, 1, 104–110.

Coltheart, M. (1978). Lexical access in simple reading tasks. In G. Underwood (Ed.), *Strategies of information processing* (pp. 151–216). New York, NY: Academic Press.

Coltheart, M. (1980). Deep dyslexia: A review of the syndrome. In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (pp. 22–47). London, UK: Routledge & Kegan Paul.

Coltheart, M., Rastle, K., Perry, C., & Langdon, R. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 208(1), 204–256.

Cortese, M. J., & Simpson, G. B. (2000). Regularity effects in word naming: What are they? *Memory & Cognition*, 28(8), 1269–1276.

Crisp, J., Howard, D., & Ralph, M. A. L. (2011). More evidence for a continuum between phonological and deep dyslexia: Novel data from three measures of direct orthography-to-phonology translation. *Aphasiology*, 25(5), 615–641.

Crisp, J., & Lambon Ralph, M. A. (2006). Unlocking the nature of the phonological–deep dyslexia continuum: The keys to reading aloud are in phonology and semantics. *Journal of Cognitive Neuroscience*, 18(3), 348–362.

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.

Dilkina, K., McClelland, J. L., & Plaut, D. C. (2008). A single-system account of semantic and lexical deficits in five semantic dementia patients. *Cognitive Neuropsychology*, 25(2), 136–164.

Ellis, A. W., & Marshall, J. (1978). Semantic errors or statistical flukes? A note on Allport's "on knowing the meaning of words we are unable to report". *Quarterly Journal of Experimental Psychology*, 30, 569–575.

Ellis, A. W., & Monaghan, J. (2002). Reply to Strain, Patterson, and Seidenberg (2002). *Journal of Experimental Psychology: Learning, Memory & Cognition*, 28(1), 215–220.

Farrar, W. T., & Van Orden, G. C. (2001). Errors as multistable response options. *Nonlinear Dynamics, Psychology, and Life Sciences*, 5(3), 223–265.

Friedman, R. B. (1996). Recovery from deep alexia to phonological alexia: Points on a continuum. *Brain and Language*, 52, 114–128.

Frost, R. (1998). Towards a strong phonological theory of visual word recognition: True issues and false trials. *Psychological Bulletin*, 123(1), 71–99.

Funnell, E. (1983). Phonological processes in reading—New evidence from acquired dyslexia. *British Journal of Psychology*, 74, 159–180.

Glosser, G., & Friedman, R. B. (1990). The continuum of deep/phonological alexia. *Cortex*, *26*, 343–359.

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.

Graham, K. S., Hodges, J. R., & Patterson, K. (1994). The relationship between comprehension and oral reading in progressive fluent aphasia. *Neuropsychologia*, *32*(3), 299–316.

Graham, N. L., Patterson, K., & Hodges, J. R. (2000). The impact of semantic memory impairment on spelling: Evidence from semantic dementia. *Neuropsychologia*, *38*(2), 143–163.

Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*(3), 491–528.

Harm, M. W., & Seidenberg, M. S. (2001). Are there orthographic impairments in phonological dyslexia? *Cognitive Neuropsychology*, *18*(1), 71–92.

Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, *111*(3), 662–720.

Hinton, G. E. (1989). Connectionist learning procedures. *Artificial Intelligence*, *40*(1–3), 185–234.

Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, *98*(1), 74–95.

Hodges, J. R., & Patterson, K. (2007). Semantic dementia: A unique clinicopathological syndrome. *Lancet Neurology*, *6*(11), 1004–1014.

Jared, D. (2002). Spelling–sound consistency and regularity effects in word naming. *Journal of Memory and Language*, *46*(4), 723–750.

Jared, D., McRae, K., & Seidenberg, M. S. (1990). The basis of consistency effects in word naming. *Journal of Memory and Language*, *29*(6), 687–715.

Jefferies, E., Ralph, M. A. L., Jones, R., Bateman, D., & Patterson, K. (2004). Surface dyslexia in semantic dementia: A comparison of the influence of consistency and regularity. *Neurocase*, *10*(4), 290–299.

Jefferies, E., Sage, K., & Lambon Ralph, M. A. (2007). Do deep dyslexia, dysphasia and dysgraphia share a common phonological impairment? *Neuropsychologia*, *45*(7), 1553–1570.

Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *PNAS*, *96*(13), 7592–7597.

Jones, G. V. (1985). Deep dyslexia, imageability, and ease of prediction. *Brain and Language*, *24*, 1–19.

Kay, J., Lesser, R., & Coltheart, M. (1992). *Psycholinguistic Assessments of Language Processing in Aphasia (PALPA)*. Hove, UK: Lawrence Erlbaum Associates.

Keidel, J. L., Welbourne, S. R., & Lambon Ralph, M. A. (2010). Solving the paradox of the equipotential and modular brain: A neurocomputational model of stroke vs. slow-growing glioma. *Neuropsychologia*, *48*(6), 1716–1724.

Kello, C. T., & Plaut, D. C. (2003). Strategic control over rate of processing in word reading: A computational investigation. *Journal of Memory and Language*, *48*(1), 207–232.

Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, *22*(1), 79–86.

Lambon Ralph, M. A., & Ehsan, S. (2006). Age of acquisition effects depend on the mapping between representations and the frequency of occurrence: Empirical and computational evidence. *Visual Cognition*, *13*, 928–948. doi:10.1080/13506280544000110

Lambon Ralph, M. A., & Graham, N. L. (2000). Acquired phonological and deep dyslexia. *Neurocase*, *6*, 141–178.

Lambon Ralph, M. A., McClelland, J. L., Patterson, K., Galton, C. J., & Hodges, J. R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *Journal of Cognitive Neuroscience*, *13*, 341–356.

Lendrem, W., & Lincoln, N. B. (1985). Spontaneous recovery of language in patients with aphasia between 4 and 34 weeks after stroke. *Journal of Neurology, Neurosurgery, and Psychiatry*, *48*(8), 743–748.

Marshall, J. C., & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research*, *2*(3), 175–199.

Marshall, J. C., & Newcombe, F. (1980). The conceptual status of deep dyslexia: An historical perspective. In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (pp. 1–20). London, UK: Routledge & Kegan Paul.

McCarthy, R., & Warrington, E. K. (1986). Phonological reading: Phenomena and paradoxes. *Cortex*, *22*, 359–380.

McCusker, L., Hillinger, M., & Bias, R. (1981). Phonological recoding and reading. *Psychological Bulletin*, *89*(2), 217–245.

Monaghan, J., & Ellis, A. W. (2002). What exactly interacts with spelling–sound consistency in word naming? *Journal of Experimental Psychology: Learning, Memory & Cognition, 28*(1), 183–206.

Mummery, C. J., Patterson, K., Price, C. J., Ashburner, J., Frackowiak, R. S. J., & Hodges, J. R. (2000). A voxel-based morphometry study of semantic dementia: Relationship between temporal lobe atrophy and semantic memory. *Annals of Neurology, 47*(1), 36–45.

Nelson, M. E., & Bower, J. M. (1990). Brain maps and parallel computers. *Trends in Neurosciences, 13*(10), 403–408.

Nolan, K. A., & Caramazza, A. (1982). Modality-independent impairments in word processing in a deep dyslexic patient. *Brain and Language, 16,* 237–264.

Patterson, K. (1982). The relation between reading and phonological coding: Further neuropsychological observations. In A. W. Ellis (Ed.), *Normality and pathology in cognitive functioning* (pp. 77–111). London, UK: Academic Press.

Patterson, K., & Hodges, J. R. (1992). Deterioration of word meaning: Implications for reading. *Neuropsychologia, 3*(12), 1024–1040.

Patterson, K., & Lambon Ralph, M. A. (1999). Selective disorders of reading? *Current Opinion in Neurobiology, 9*(2), 235–239.

Patterson, K., Lambon Ralph, M. A., Jefferies, E., Woollams, A., Jones, R., Hodges, J. R., et al. (2006). Pre-semantic cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience, 18*(2), 169–183.

Patterson, K., & Marcel, A. J. (1992). Phonological ALEXIA or PHONOLOGICAL alexia? In J. Alegria, J. Holender, J. Junça de Morais, & M. Radeau (Eds.), *Analytic approaches to human cognition* (pp. 259–274). Amsterdam, The Netherlands: Elsevier Science Publishers.

Patterson, K., Plaut, D. C., Seidenberg, M. S., Behrmann, M., & Hodges, J. R. (1996). Connections and disconnections: A connectionist account of surface dyslexia. In J. A. Reggia, E. Ruppin, & R. S. Berndt (Eds.), *Neural modeling of brain and cognitive disorders* (pp. 177–199). New York, NY: World Scientific.

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* (pp. 131–181). London, UK: Oxford University Press.

Patterson, K., Suzuki, T., & Wydell, T. (1996). Interpreting a case of Japanese phonological alexia: The key is in phonology. *Cognitive Neuropsychology, 13*(6), 803–822.

Patterson, K., Suzuki, T., Wydell, T., & Sasanuma, S. (1995). Progressive aphasia and surface alexia in Japanese. *Neurocase, 1*(2), 155–165.

Pedersen, P. M., Jørgensen, H. S., Nakayama, H., Raaschou, H. O., & Olsen, T. S. (1995). Aphasia in acute stroke: Incidence, determinants, and recovery. *Annals of Neurology, 38*(4), 659–666.

Perry, C., Ziegler, J. C., & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review, 114*(2), 273–315.

Plaut, D. C. (1995a, June). *Connectionist modeling of the breakdown and recovery of reading via meaning*. Paper presented at Neural Modeling of Brain and Cognitive Disorders, College Park, MD.

Plaut, D. C. (1995b). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology, 17*(2), 291–321.

Plaut, D. C. (1996). Relearning after damage in connectionist networks: Toward a theory of rehabilitation. *Brain and Language, 52*(1), 25–82.

Plaut, D. C. (1996). Structure and function in the lexical system: Insights from distributed models of word reading and lexical decision. *Language and Cognitive Processes, 12,* 767–808.

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., & Kello, C. T. (1998). The emergence of phonology from the interplay of speech comprehension and production: A distributed connectionist approach. In M. B. (Ed.), *The emergence of language*. Mahwah, NJ: Lawrence Erlbaum Associates.

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.

Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology, 10*(5), 377–500.

Rapcsak, S. Z., Beeson, P. M., Henryc, M. L., Leyden, A., Kim, E., Rising, K., et al. (2009). Phonological

dyslexia and dysgraphia: Cognitive mechanisms and neural substrates. *Cortex*, *45*(5), 575–591.

Reynolds, M., & Besner, D. (2005). Basic processes in reading: A critical review of pseudohomophone effects in reading aloud and a new computational account [Review]. *Psychonomic Bulletin & Review*, *12*(4), 622–646.

Rohde, D. L. T. (2000). Lens the light efficient neural network simulator v2.63. Retrieved from http://tedlab.mit.edu/~dr/Lens/

Schwartz, M. F., Marin, O. S. M., & Saffran, E. M. (1979). Dissociations of language function in dementia—Case-study. *Brain and Language*, *7*(3), 277–306.

Seidenberg, M. S. (1995). Visual word recognition: An overview. In J. L. Miller & P. D. Eimas (Eds.), *Speech, language, and communication* (2nd ed., pp. 138–172). San Diego, CA: Academic Press.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523–568.

Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.

Shallice, T., & Warrington, E. K. (1980). Single and multiple component central dyslexic syndromes. In M. Coltheart, K. Patterson, & J. Marshall (Eds.), *Deep dyslexia* (pp. 119–145). London, UK: Routledge & Kegan Paul.

Shibahara, N., Zorzi, M., Hill, M. P., Wydell, T., & Butterworth, B. (2003). Semantic effects in word naming: Evidence from English and Japanese Kanji. *Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology*, *56*(2), 263–286.

Smith, F. (1973). *Pscholinguistics and reading*. New York, NY: Holt, Rinehart and Winston.

Strain, E., Patterson, K., & Seidenberg, M. S. (1995). Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *21*(5), 1140–1154.

Strain, E., Patterson, K., & Seidenberg, M. S. (2002). Theories of word naming interact with spelling–sound consistency. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *28*(1), 207–214.

Taraban, R., & McClelland, J. L. (1987). Conspiracy effects in word pronunciation. *Journal of Memory and Language*, *26*, 608–631.

Thomas, M. S. C., Grant, J., Barham, Z., Gsodl, M., Laing, E., Lakusta, L., et al. (2001). Past tense formation in Williams syndrome. *Language and Cognitive Processes*, *16*(2–3), 143–176.

Tree, J. J., & Kay, J. (2006). Phonological dyslexia and phonological impairment: An exception to the rule? *Neuropsychologia*, *44*(14), 2861–2873.

Weekes, B., & Coltheart, M. (1996). Surface dyslexia and surface dysgraphia: Treatment studies and their theoretical implications. *Cognitive Neuropsychology*, *13*(2), 277–315.

Welbourne, S. R., & Lambon Ralph, M. A. (2005a). 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.

Welbourne, S. R., & Lambon Ralph, M. A. (2005b). Using computational, parallel distributed processing networks to model rehabilitation in patients with acquired dyslexia: An initial investigation. *Aphasiology*, *19*(9), 789–806.

Welbourne, S. R., & Lambon Ralph, M. A. (2006). Phonological and surface dyslexia in a single PDP model of reading, *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006* (pp. 2359–2364). Vancouver, Canada.

Welbourne, S. R., & Lambon Ralph, M. A. (2007). Using PDP models to simulate phonological dyslexia: The key role of plasticity-related recovery. *Journal of Cognitive Neuroscience*, *19*, 1125–1139.

Wilson, S. M., Brambati, S. M., Henry, R. G., Handwerker, D. A., Agosta, F., Miller, B. L., et al. (2009). The neural basis of surface dyslexia in semantic dementia. *Brain*, *132*(1), 71–86.

Woollams, A., Lambon Ralph, M. A., Plaut, D. C., & Patterson, K. (2007). SD-squared: On the association between semantic dementia and surface dyslexia. *Psychological Review*, *114*(2), 316–339.

Young, M. P., Scannell, J. W., & Burns, G. (1995). *The analysis of cortical connectivity*. Heidelberg, Germany: Springer-Verlag.

Zevin, J. D., & Seidenberg, M. S. (2006). Simulating consistency effects and individual differences in nonword naming: A comparison of current models. *Journal of Memory and Language*, *54*(2), 145–160.

Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(4), 1131–1161.