ELSEVIER

Cognition 74 (2000) 91–113

COGNITION

www.elsevier.com/locate/cognit

Attractor dynamics in word recognition: converging evidence from errors by normal subjects, dyslexic patients and a connectionist model

Peter McLeod[a,*], Tim Shallice[b], David C. Plaut[c]

[a]Department of Experimental Psychology, Oxford University, Oxford, UK
[b]Institute of Cognitive Neuroscience, University College London, London, UK
[c]Psychology Department, Carnegie Mellon University, Pittsburgh, PA, USA

Received 21 October 1998; received in revised form 22 June 1999; accepted 24 September 1999

Abstract

People make both semantic and visual errors when trying to recognise the meaning of degraded words. This result mirrors the finding that deep dyslexic patients make both semantic and visual errors when reading aloud. We link the results with the demonstration that a recurrent connectionist network which produces the meaning of words in response to their spelling pattern produces this distinctive combination of errors both when its input is degraded and when it is lesioned. The reason why the network can simulate the errors of both normal subjects and patients lies in the nature of the attractors which it develops as it learns to map orthography to semantics. The key role of attractor structure in the successful simulation suggests that the normal adult semantic reading route may involve attractor dynamics.

Keywords: Attractor dynamics; Connectionist modelling; Word recognition; Word recognition errors

1. Introduction

Neuropsychology contains many ‘signature’ phenomena, surprising patterns of error that occur after an information processing mechanism is damaged. If the same

* Corresponding author. Fax: +44-1865-790-819.
E-mail addresses: peter.mcleod@psy.ox.ac.uk (P. McLeod); t.shallice@ucl.ac.uk (T. Shallice); plaut@cmu.edu (D.C. Plaut)

0010-0277/00/$ - see front matter
PII: S0010-0277(99)00067-0

pattern occurs when normal subjects process degraded stimuli then it is natural to conclude that both patient and normal data reflect malfunctioning of the same mechanism. This conclusion would be strengthened if the two patterns could be linked by a common model in which lesioning (simulating the patient data) produced qualitatively the same effects as degrading input to it (simulating the normal data). If the computational model exhibits other qualitative characteristics broadly in line with normal performance, then such a convergence provides evidence that the cognitive system contains a mechanism which operates by a related computational algorithm.

This paper demonstrates just such a link between errors in the identification of written words by normal subjects, neurological patients and a connectionist model. This type of three-way link was first established by Mozer, Behrmann, and colleagues in the context of neglect dyslexia. Lateralized damage to the attentional mechanism of MORSEL, a connectionist model of word and object recognition (Mozer, 1991) reproduced the pattern of word reading errors exhibited by neglect dyslexic patients (Mozer & Behrmann, 1990), and a lateral bias on attention in the normal model reproduced analogous effects on the latencies of normal subjects when their attention was similarly biased (by underlining the ends of letter strings; Behrmann, Moscovitch & Mozer, 1991). The current work establishes an analogous convergence in the context of deep dyslexia. We begin by briefly reviewing earlier work which established that lesions to a connectionist model which maps spelling patterns onto semantics can account for both the semantic errors which are the striking aspect of reading by deep dyslexic patients and the visual errors which, surprisingly, accompany them. The key to the network's ability to simulate the performance of patients is the structure of the attractors which it acquires as it learns to map orthography to semantics. We then report experiments in which normal subjects make both semantic and visual errors when searching for semantically defined target words under degraded viewing conditions. In the third section we demonstrate that the model which mimicked the patient data produces the same effects if, instead of being lesioned, its input is degraded by being masked before a response has been produced. In the final section we show that the similarities of the effects of damage and the effects of input degradation to the model are a result of the computational principles of attractor networks. This leads to the conclusion that the mechanism which identifies words via their meaning in people has attractor dynamics.

2. The co-occurrence of semantic and visual reading errors in deep dyslexia

Deep dyslexia is an acquired disorder in which patients make semantic errors in oral reading. For example, a patient shown the word ILL read it as 'sick' (Marshall & Newcombe, 1966). Explanations of the syndrome are generally cast within models of the reading system in which phonological and semantic routes from print to sound can be distinguished. It is assumed that the phonological route(s) used in normal oral reading are inoperative in the patients and they must rely on the semantic route. In addition, the semantic route is generally assumed to be impaired. Deep dyslexics vary in where the impairment lies in the semantic route. In the 'input' variety reading comprehension is inferior to auditory comprehension which points to the locus of damage being prior to the semantic system (see Friedman & Perlman, 1982; Shallice & Warrington, 1980).

In the literature on acquired dyslexia since 1973 over 50 patients have been described who make semantic errors in reading. While the existence of semantic errors is the cardinal symptom of deep dyslexia, it has been shown to occur with a set of other symptoms. The first is that these patients also make visual errors (e.g. SCANDAL read as 'sandals')[1]. A variety of other symptoms have been described: greater confidence in visual errors than in semantic errors; 'surprisingly good' lexical decision (Coltheart, 1980); the existence of visual-then-semantic errors (e.g. SYMPATHY read as 'orchestra'); better performance reading concrete (imageable) than abstract words. (For a general discussion of deep dyslexia see Coltheart, Patterson & Marshall, 1987; Shallice, 1988; Van Orden, 1987).

A theoretical explanation for the co-occurrence of these symptoms was provided by Hinton and Shallice (1991) and by Plaut and Shallice (1993a) who investigated the effects of damage to connectionist networks that mapped orthography to semantics on the basis of the following computational principles:

(a) orthographic and semantic representations are distributed over separate groups of units such that similar patterns represent similar words in each domain, but similarity is unrelated between domains;
(b) Learning operates by adjusting weights on connections between units so as to perform gradient descent in a measure of performance on the task of mapping orthography to semantics;
(c) The process of mapping orthography to semantics in the trained system is accomplished through the operation of *attractors*. In an attractor network, the weights on connections between units cause particular patterns of activity to be stable. When the network is placed in an initial state by an input the interactions among units will cause the network to settle, over time, into the nearest stable state. This state corresponds to its interpretation of the input.

The primary discovery by Hinton and Shallice was that, wherever the networks were lesioned, the result was the co-occurrence of semantic errors (e.g. the input string CAT producing the output 'dog') and visual errors (e.g. the input string LOG producing the output 'dog'). That is, the lesioned network produced the same pattern of errors as found in deep dyslexia[2]. Plaut and Shallice (1993a) went on to establish the generality of these findings across different network architectures and learning

[1] Three patients have been described by Caramazza and Hillis (1990); Hillis, Rapp Romani and Caramazza (1990) who make semantic errors without visual errors. For discussion of these exceptional cases see Plaut and Shallice (1993a).

[2] Newton and Barry (1997) have argued that in deep dyslexia visual errors can arise from an unrelated visual deficit. But this is implausible in input deep dyslexics, where the errors occur more frequently on words in the semantic categories which are *most* difficult for the patient to read (see Shallice &Warrington, 1980); in other words the impairments at semantic and visual levels interact.

algorithms and to demonstrate that the full range of symptoms of deep dyslexic patients described above arise when attractor networks are lesioned. But the fundamental finding from lesioning the attractor networks was the ubiquitous co-occurrence of visual errors and semantic errors following damage. The use of attractors for performing the task was critical for obtaining the pattern of errors which mimicked that produced by the patients. These errors occurred because damage to the network altered the layout of the attractors in such a way that an input might be captured by the attractor for a related word.

This co-occurrence is a consequence of the necessity to overcome the fundamental bias in connectionist networks to give similar responses to similar inputs. This bias supports effective generalization in most domains (e.g. mapping orthography to phonology in oral reading; see Plaut, McClelland, Seidenberg and Patterson (1996)), but is problematic when mapping between arbitrarily related domains like orthography and semantics. Apart from morphological variants, words that look the same do not tend to have similar meanings. Developing attractors for word meanings can help a network overcome the detrimental effects of a bias based towards similarity. Specifically, in an attractor network, visually similar words are free to generate similar initial semantic patterns as long as these patterns fall within different basins of attraction. The attractors serve to pull these initially similar patterns apart into the appropriate distinct final patterns. One consequence of this solution, however, is that there is are regions in semantic space where neighbouring attractors correspond to *visually* similar words.[3] Damage may distort these basins, occasionally causing the normal initial semantic pattern of a word to be captured within the basin of a visually similar word. In other regions of semantic space, neighbouring attractors correspond to semantically related words, and damage may cause initial patterns of activation in these regions to give rise to semantic errors. Essentially, the layout of attractor basins must be sensitive to both visual and semantic similarity, and so these metrics are reflected in the types of errors that occur as a result of damage.

If human word reading via meaning operates via attractors that are sensitive to both visual and semantic similarity it is possible that, with appropriate presentation conditions, normal skilled readers would, like deep dyslexic patients, exhibit visual

[3] It might seem counter intuitive that the attractor basins for two visually similar words can be nearby in semantic space even though the words are semantically unrelated. The reason this is possible stems from the geometric properties of high dimensional spaces (see Hinton and Shallice (1991), Appendix A). The attractor patterns for two semantically unrelated words can be thought of as two random points in semantic space. The natural region in this space for the attractor basins of the two words to be close together is about halfway between the attractor points for the words. If this midpoint were closer to the attractor for another word (and, hence, likely to fall within its basin) the network would have difficulty avoiding errors. Hinton and Shallice showed, however, that the likelihood that the midpoint between two random points in a space is closer to any of a large number of other random points than it is to the original points becomes vanishingly small as the dimensionality of the space increases. Thus it is relatively easy for the network to shape the attractor basins of visually similar words so that they are nearby in semantic space without the basins for other words getting in the way and causing errors.

and semantic effects in their reading behavior. The following experiments explore this possibility.

3. The co-occurrence of semantic and visual errors by normal subjects following rapid serial visual presentation of words

If a similar pattern of errors were shown by normal subjects with degraded input as were made by patients the assumption might be made that the errors stemmed from malfunction of the same mechanism, in one case as a result of damaging it, in the other the result of giving it inadequate input. The predictable result that degraded visual input leads to visual errors has been observed many times (e.g. Morton, 1969). The suggestion that normal subjects reading brief masked words might, like deep dyslexic patients, also produce errors which were semantically related to the stimuli was first made by Allport (1977). However, the conventional masked words paradigm is not ideal for the demonstration of semantic errors. A trial typically consists of a single word followed by a mask. Any semantic errors which are generated are unlikely to be orthographically or phonologically related to the stimulus. So potential semantic paralexias might be rejected by a postperceptual editing process because they did not fit some of the letters which had been detected. To investigate the production of semantic errors following brief masked presentation it is necessary to use an experimental paradigm where the role of recognised letters in checking candidate responses is minimized.

We required subjects to perform a semantic search among words presented under conditions of rapid serial visual presentation (RSVP). In this technique a string of words is presented successively, at the same spatial location, each for too short a time to be clearly identified. The subject must report whether any word satisfies a pre-specified semantic criterion. Phenomenologically the presentation conditions create a dense, confusable, and continuously changing set of visual traces at the letter level. Although the viewer gets a jumbled impression of both letters and words, it is difficult to use letter-level information to guide or monitor word-level report because it is difficult to tell when the various traces occurred. For example, the belief that you had seen the word DOG would not necessarily be ruled out by the belief that there was an initial C because that might have occurred in the previous or subsequent letter string.

Since we wished subjects to access semantics with visual information (to mimic the use of the semantic route by deep dyslexic patients) we employed a procedure which should reduce the possibility of phonological semantic access (see Van Orden, 1987). We used a rate of presentation faster than standard speech input and we included polysyllabic words since rapid online phonological accessing of semantics using all phonemes in parallel seems likely to be limited to single syllable words. Finally, the use of a paradigm in which targets are separated from the background words in terms of semantic fields rather than on the basis of phonology makes the Orthographic → Semantic mapping for the target less subject to interference than the Orthographic → Phonological → Semantic mapping.

3.1. Experiment 1

3.1.1. Method

3.1.1.1. Subjects These were 20 people aged between 20 and 40 years selected from the subject panel of the MRC Applied Psychology Unit.

3.1.1.2. Stimuli Subjects viewed a list of words presented under RSVP conditions. Their task was to identify any word satisfying the semantic criterion 'things not normally found inside a house'. Words meeting this criterion were selected from a wide range of semantic categories. For example, targets included *aeroplane*, *manure*, *wolf*, *pier* and *meadow*. This criterion produces a very wide field of potential targets, so an error produced at random is unlikely to be semantically related to the target. This makes the detection of real semantic errors possible. The non-target words in the list were the names of things which would commonly be found in a house. They were roughly matched for word length (four to nine letters) with the targets. They also fell in many different semantic categories. For example, background items included *cat*, *chocolate*, *scissors*, *earring*, and *parsley*.

Each list contained between 15 and 25 words. There were either one or two targets in each list. A target never appeared in the first five or last five words in the list but was equally likely to appear in any other position. There was a pool of 240 background words. A random selection from these was made up on each trial to provide the required number of filler items. Background words were repeated at random across trials, but never occurred twice in the same list.

3.1.1.3. Procedure Prior to the experimental session subjects were given practice at the task. Subjects received 12 lists with the interword interval (IWI) gradually decreasing from 1 s to 160 ms. This was followed by 16 practice trials at 160 ms. The experimental session consisted of 100 lists with an interword interval of 160 ms.

List presentation lasted 3–4 s followed by an intertrial interval of about 10 s. Subjects reported any target or partial information about possible targets verbally to the experimenter. They were given no feedback about the accuracy of their reports. The words were displayed on an Electronic Visuals oscilloscope driven by a Cambridge Electronic Design computer. They were composed of upper case letters generated from a 9×5 dot matrix. The system took approximately 1 ms to display a five-letter word. The oscilloscope display was refreshed at 20 ms intervals. Subjects were seated roughly 80 cm from the display: a five-letter word occupied about 3° of visual angle. The background illumination of the room was very dim: the characters appeared bright against a blank screen.

3.1.2. Results and discussion

Sixty-one percent of targets were correctly identified. To the remaining targets subjects either made no response or they reported a word which had not been presented. The errors could reflect the pickup of visual information from background words (or the target), or semantic information from the target word. The aim of the initial analysis was to separate those errors which appeared to be influenced primarily by visual information before testing the remaining errors for evidence of the influence of semantic information derived from the target word.

An error response was classified as a candidate Visual error if it had only one letter different from any background item on the trial (e.g. 'bomb' given as a response to a list containing the word COMB), or if it contained a block of four or more consecutive letters in common with any background item (e.g. 'antelope' given as a response to a list containing ENVELOPE). An error response was classed as a candidate Mixed Visual and Semantic error if it had three or more consecutive letters in common with the target or differed overall by no more that one letter. (Different criteria were used to take into account the differing baseline probabilities of a response being visually similar to any item in the list compared to its being similar specifically to the target.)

Fifty-five error responses were produced. Of these four were eliminated because they were not outdoor objects, 15 were classified as Visual errors and two were candidates for Mixed Visual and Semantic errors. This left 34 errors without a strong visual relation to any word in the list for the semantic relatedness analysis.

3.1.2.1. Semantic errors The non-visual errors and the targets on the trials when they were produced are given in the Appendix A. Inspection of these shows many immediately striking examples which suggest semantic links between target and response: DEER produced the response 'stag'; CLOUD produced 'helicopter'; SNOW produced 'to do with storms, weather' and so on. Convincing as such examples may seem, the probability of two words drawn at random from the English language appearing to have a semantic relation is not negligible (see Ellis & Marshall, 1978; Williams & Parkin, 1980). It is therefore necessary to establish what the chance rate for semantic relatedness would be between our stimuli and the putative semantic error responses, and to test the rate found in the experiment against this.

Each error produced in the experiment was paired with the target(s) which appeared on that trial (e.g. deer – *stag*) and was also randomly paired with one of the other targets which had produced an error (e.g. cloud – *stag*). The lists of target + error pairs were given to judges who were given a rough outline of the experiment. They were told that half the word pairs on the list were errors generated in the experiment and half were random reassignments of errors to stimuli. Their task was to classify the semantic relatedness of each pair into one of the four categories: Very Close, Close, Possible or None. The classification pattern of the randomly assigned pairs gives a measure of the degree of semantic relatedness which might occur by chance. This pattern can be compared with that obtained from pairings produced by subjects in the experiment. Table 1 shows the ratings pooled across four judges for the semantic relatedness of the real target + error pairs and the random stimulus response pairings. Semantic relation between target and error was far more likely for pairs obtained in the experiment than would be expected by chance ($\chi^2 = 40.1$, $P < 0.001$). (This value was calculated pooling the scores for Very Close, Close and Possible as evidence of semantic relatedness, and contrasting

Table 1
The probability of a target-response pair being rated as having a given degree of semantic relatedness in Experiment 1 (Section 3.1). Results are given for experimentally obtained responses and for random re-pairing of these with the targets

	Experiment	Random re-pairing
Very close	0.29	0.05
Close	0.21	0.04
Possible	0.26	0.18
Not	0.24	0.72

this with Not. The same conclusion is reached if the division is made between Very Close and Close versus Possible and Not.) We may conclude that people produce error responses which are semantically related to the targets when they try to identify the meaning of words under conditions of Rapid Serial Visual Presentation. This echoes the finding by Potter and Lombardi (1990) that words which have been primed by prior presentation may be reported when a semantically related word is presented in an RSVP string.

3.1.2.2. Visual errors Fifteen putative visual errors were observed. Despite the strict criterion - four or more consecutive letters in common between stimulus and response or only one letter different – it is possible that they might arise from a response being visually similar to one of the words in a list by chance rather than from a visual misidentification process. The procedure of Ellis and Marshall (1978) of random reassignment of error responses to stimuli was followed, and then visual similarity was assessed with the criteria previously adopted for visual similarity to background items. Random reassignment of the errors pairs produced four pseudovisual errors by comparison with the 15 errors generated in the experiment. Real Visual errors are significantly more frequent than pseudovisual errors (McNemar Test, $P < 0.001$). They are therefore a real phenomenon not an artifactual effect of chance association.

3.1.2.3. Mixed errors Too few putative Mixed Visual and Semantic errors occurred to allow analysis of this error type.

3.1.3. Conclusion

Subjects were exposed to a rapidly changing string of words and tried to report any which fell into a particular semantic category. With a display in which it was difficult to use letter-level information to check responses, they produced both semantic errors and visual errors. That is, they reported seeing words which were not there, but which were related either semantically or visually to those that were. Thus, normal subjects produced an error pattern with the degraded input conditions produced by RSVP that is qualitatively similar to that produced by deep dyslexic patients.

3.2. Experiment 2

In Experiment 1 (Section 3.1) the inter-word interval was set at a level where pilot experiments suggested people were just beginning to be unable to perform the task of search for semantically defined words under RSVP conditions with certainty. In Experiment 2 we explored the generality of this result by testing whether the phenomenon revealed in Experiment 1 (Section 3.1), the co-occurrence of visual and semantic errors, would continue to occur across a range of stimulus durations, down to a point where the task could hardly be performed at all.

3.2.1. Method

3.2.1.1. Subjects There were 15 subjects taken from the Oxford University Psychology Department panel, age range 20–48.

3.2.1.2. Stimuli These were the same as those used in Experiment 1 (Section 3.1.2). There was only one target per list.

3.2.1.3. Procedure The procedure was the same as Experiment 1 (Section 3.1.3) except that in the practice phase the inter-word interval was reduced in steps of 100 ms from 1000 to 300 ms and then in steps of 20 to 100 ms. The experimenter then calculated the stimulus duration at which an individual subject started to miss targets. If this was T ms, the subjects was then run with inter-word intervals of (T-40), 0.75(T-40) and 0.5(T-40) ms. The three conditions produced median inter-word intervals across subjects of 156, 112 and 75 ms respectively. Then they were shown 20 lists at each of the three durations. The order of presentation of the three conditions was randomised across subjects.

3.2.2. Results

Table 2 shows overall performance at the different stimulus durations. Naturally the probability of detecting the target falls as the stimulus duration is reduced, although even at a median presentation rate of about 12 words per second 20% of the targets were identified. As performance fell both errors and omissions (trials on which the subjects made no response) increased by a roughly comparable amount. There were 27 errors at 156 ms, 49 at 112 ms, and 92 at 75 ms. The increased error rate at shorter stimulus durations is reliable across subjects by the Wilcoxon test

Table 2
Performance in Experiment 2 (Section 3.2) as a function of stimulus duration

Stimulus duration (ms)	*P*(correct)	*P*(omission)	*P*(error)
75	0.23	0.44	0.33
112	0.45	0.37	0.18
156	0.60	0.30	0.10

(P(error) at 112 ms $> P$(error) at 156 ms, T = 15.5, $P < 0.01$; P(error) at 75 ms $>$ P(error) at 112 ms, W = 15, $P < 0.01$).

The increase in error rate as stimulus duration is reduced is, on some theories, surprising. On one traditional view of information processing, as represented by the logogen model (Morton, 1969) for example, stimulus information integrates over time until a threshold is exceeded. A response is then produced. When the stimulus duration is reduced to one at which there is insufficient evidence to produce a response on every trial the expected result would be fewer responses of any sort. That is, as stimulus duration reduces it would be expected that omissions would rise and responses, whether correct or erroneous, would fall.

3.2.2.1. Visual errors The same procedure was used for classifying errors as Visual (i.e. four letters in common with a background or target word, or only one letter different) as in Experiment 1 (Section 3.1). This accounted for 9, 12 and 14 of the errors at 156, 112 and 75 ms, respectively. In each case the number of visual errors found in the experiment was greater than expected by chance using the rate of production of pseudovisual errors produced with the same stimuli and background items as in Experiment 1 (2, 3.4 and 5, respectively) ($\chi^2 = 26.4$, 23.4 and 17.4. respectively, $P < 0.001$ in each case).

3.2.2.2. Semantic errors This left 19, 37 and 78 errors at 156, 112 and 75 ms. As before there were a number of striking errors, suggesting an influence of the semantic field of the target word on the response: CHURCH produced the response ‘steeple’, SAND produced ‘beach’, SPIRE produced ‘Nelson’s column’. The errors were assessed for semantic relatedness to the target word using the same 4-point scale as in Experiment 1 (Section 3.1). Table 3 shows the proportion of errors falling in each of the semantic relatedness classes, for the three judges. Using the assessments of the random pairings in Experiment 1 (Section 3.1) as the basis for the number expected in each category by chance with this stimulus set there was positive evidence of semantic influence of target on response at 75 ms ($\chi^2 = 13.3$, $P < 0.001$) and at 112 ms ($\chi^2 = 16.6$, $P < 0.001$) but not at 156 ms ($\chi^2 = 2.19$, n.s.)[4]. The failure to find a semantic relatedness effect in the condition which is closest to Experiment 1 (Section 3.1) in terms of stimulus duration and percent

Table 3
The probability of a target-response pair being rated as having a given degree of semantic relatedness in Experiment 2 (Section 3.2) at three stimulus durations

	Stimulus duration		
	156 ms	112 ms	75 ms
Very close	0.00	0.05	0.05
Close	0.00	0.22	0.14
Possible	0.42	0.30	0.29
Not	0.58	0.43	0.52

correct performance is surprising. However, since Experiment 1 (Section 3.1) was based on a sample six times the size of the one here we will assume that the result of Experiment 1 (Section 3.1) is correct.[5]

3.2.3. Conclusions

The three experiments with normal subjects show that over a range of stimulus durations which give a probability of target detection from about 0.6 to about 0.25 the probability of both omissions and errors increases as the probability of target detection decreases. Across this range of stimulus duration the errors which subjects produce are related to both the visual and semantic characteristics of stimulus words.

4. The co-occurrence of semantic and visual errors in a connectionist model following brief masked presentation

The co-occurrence of semantic and visual errors in deep dyslexic patients and in normal readers trying to identify the meaning of words under degraded viewing conditions supports the view that the data from patients reflect the malfunction of a mechanism used in normal reading. Since the patient data has been mimicked successfully with a model in which orthography is mapped to semantics via attractors (Plaut & Shallice, 1993a) it might seem appropriate to model the semantic route in normal reading in the same way. It would then become necessary to demonstrate that a network which exhibits the deep dyslexic error pattern when damaged will also produce the co-occurrence of error types in normal operation when inputs are presented too briefly to be processed accurately.

Ideally, the network would perform the same task on stimuli presented as they were to the human subjects. However, the networks trained by Plaut and Shallice (1993a) can only process one word at a time so a direct analogue of the RSVP task is not possible. Also the network does not have the attentional and control mechanisms necessary for maintaining an *ad hoc* semantic category to define targets ('things not found indoors'). However, we assume that the main limitation to identifying words under RSVP conditions is that the time available to process each item is insufficient to guarantee correct identification, and that the subsequent word serves as a visual mask for the previous word. So we used a mode of stimulus presentation that

---

[4] These values were calculated pooling the scores for Very Close, Close and Possible as evidence of semantic relatedness, and contrasting these with Not. The same conclusions are reached if the division is made between Very Close and Close versus Possible and Not.

[5] To check this conclusion we ran a replication of Experiment 1 (Section 3.1) (i.e. an experiment with a fixed IWI of 160 ms) with a new group of 16 subjects and some minor changes in procedure. There was an overall probability of target detection of 0.48. There were 96 errors of which 52 were classified as Visual and four were eliminated as they were not in the appropriate semantic class. The remaining 40 were tested for semantic relatedness to the target using the same technique as in Experiment 1 (Section 3.1). The number of semantically related responses produced in the experiment was significantly greater than would be expected by chance ($\chi^2 = 17.4$, $P < 0.001$).

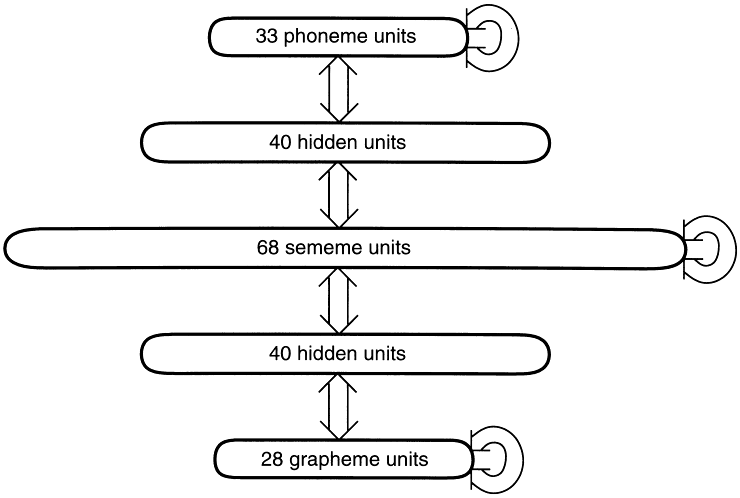

Fig. 1. The network used in the simulations. Ovals represent groups of units, and arrows represent full bidirectional connectivity within or between these groups.

approximated the effects of RSVP on individual stimulus items – naming under brief masked exposure, where the mask consists random combinations of letters which can be interpreted as a proxy for the integrated orthographic representations of non-target words.

4.1. Method

4.1.1. Network and task

The network mapped from orthography to semantics to phonology (see Fig. 1). As there is no direct mapping between orthography and phonology, the model represents reading by the semantic route alone. Sets of 40 hidden units were present between orthography (28 grapheme units) and semantics (68 sememe units) and between semantics and phonology (33 phoneme units). There was full bidirectional connectivity between adjacent layers of units, as well as between units within the grapheme, sememe, and phoneme layers. During processing, each unit adopted a real-valued activation between -1 and $+1$ which was a monotonic, non-linear (logistic) function of its summed weighted inputs from other units. The network, input and output representations, learning algorithm and training procedure were the same as employed in the Deterministic Boltzmann Machine simulations of Plaut and Shallice (1993a).[6] The Deterministic Boltzmann Machine was used because this has a fine time grain so it is possible to adjust mask and target durations precisely to produce any desired level of performance.

The network was taught a set of 40 words, eight each in the semantic categories *indoor objects*, *animals*, *body parts*, *foods*, and *outdoor objects*. All words contained

[6] They are described in detail in Section 4.1 and Appendix 1 of that paper.

Table 4
Performance of the network as a function of stimulus duration

Stimulus duration	*P*(correct)	*P*(omission)	*P*(error)
5	0.02	0.43	0.55
15	0.13	0.46	0.41
30	0.39	0.31	0.30
50	0.64	0.19	0.17
70	0.90	0.04	0.06

three or four letters, with a restricted set of letters being used in each letter position. This limited the number of alternative letters in the four possible letter positions to a total of 28. Semantics consisted of 68 features intuitively generated to fit the semantic characteristics of the 40 words. Each word was represented by a positive value on a subset of these features. On average a word had 15 positive features out of the possible 68. The phonological representation of a word was specified in terms of 33 position-specific phoneme units.[7]

4.1.2. Stimulus presentation

To implement brief masked exposure, a word was presented to the network for a fixed number of iterations and then replaced by a random mask for ten iterations. In the mask each grapheme had a probability of 0.12 of being active. This was equal to the probability of a grapheme being active averaged across the whole stimulus set. After the mask was removed all units in the network continued to update their states until the settling process had reached asymptote.

4.2. Results

Table 4 shows the performance of the network, with the probabilities of correct response, omissions and errors at different numbers of iterations before the mask was applied, corresponding to successively longer stimulus durations. Each set is the average result across 50 runs using a different random mask, where each word in the stimulus set was presented once per run. The *P*(correct) column shows how well the network performed. Naturally the longer it was allowed to settle before the mask was clamped to the input in place of the stimulus word the better was its performance. The output was categorized as an Omission if for any phoneme position no phoneme was activated above the threshold. The output was classed as an Error when one phoneme was activated in each position but at least one of these did not correspond to those in the correct response.[8]

[7] The words, semantic features and phonological representations are the same as those used by and listed in Plaut and Shallice (1993a).

[8] Averaged across stimulus duration and error type the error response is also the best match at the semantic level on 99% of trials. Thus the results are not dependent on using ‘reading aloud’ as the behavioural task for the network.

4.2.1. Error classification

When the network produced a error it was almost always one of the words it learned during the training phase (the maximum rate of responses which were not in the training set was 0.2%). Following the system used by Plaut and Shallice (1993a) the errors were categorized into four types.

- Visual: responses that were visually (but not semantically) similar to the stimulus (e.g. CAT → 'cot'). They had at least one letter in the same position as in the stimulus but the response came from a different semantic category from the stimulus.
- Semantic: responses that were semantically (but not visually) similar to the stimulus (e.g. CAT → 'dog'). They were words from the same semantic category as the stimulus, without any letter in the same position in stimulus and response.
- Mixed: responses that were both visually and semantically similar to the stimulus (e.g. CAT → 'rat'). They were words from the same semantic category as the stimulus with at least one letter in the same position.
- Unrelated: responses which came from the stimulus set but fell into none of the three categories above (e.g. CAT → 'dune').

4.2.2. Comparison with chance error rate

The frequencies of error types are shown in Table 5 for five different stimulus durations. If the model were to produce responses from the training set at random they would fall into one of the error categories. The base rate for error types which would be found if the model produced random responses was computed by pairing each word in the training set (treated as a stimulus) with each other word in the training set (treated as a response), and seeing how many of these random pairings fell into each error category. According to this calculation, the chance rates of each error type are: Visual, 30%; Semantic, 12%; Mixed, 6%; Unrelated, 52%. After five iterations the stimulus has had very little effect on the network and the proportions of errors are those which would be expected by chance – Visual 29%; Semantic 14%; Mixed 7%; Unrelated 49%. By 15 iterations the stimulus has started to influence the

Table 5
Performance of the network when presented with a brief input followed by a mask. The table shows the percentage of correct responses at each stimulus duration. The errors are shown as the percentage of those occurring which fell into each classification

Stimulus duration (iterations)	% Errors by type			
	Visual	Semantic	Mixed	Unrelated
5	29	14	7	49
15	26	30	15	31
30	12	50	18	20
50	10	67	14	8
70	10	70	15	3

outcome of the network before the imposition of the mask, and the pattern of errors has begun to deviate from chance.

The question is whether the nature of the stimulus – its orthographic structure and the semantic class it came from – influenced the error responses produced during the simulation. Were the error types which involve a relationship between stimulus and response (Visual, Semantic or Mixed) produced more frequently than would be expected given the base rate for their occurrence? This can be determined by taking the ratio of the rate of occurrence of a particular error type to the rate of Unrelated errors and comparing the ratio produced in the simulation with that expected by chance. Thus, for example, in the 50 iteration condition, Semantic errors are 8.4 times more frequent then Unrelated errors (i.e. 67%/8%). Their expected ratio by chance is 0.23 (12%/52%). Clearly, in this condition, Semantic errors are more frequent than would be expected by chance.

The ratio of Semantic errors to Unrelated errors was 0.97, 2.5, 8.4 and 23.3 for exposure durations of 15, 30, 50 and 70 iterations, respectively. In each case these values are reliably higher than the value of 0.23 which would be expected if Semantic errors occurred by chance ($\chi^2 \geq 260$; $P < 0.001$ in each case).

Visual errors would be expected to occur by chance on 0.57 (30/52) times as many occasions as Unrelated errors. They occurred on 0.84, 0.6, 1.3 and 3.3 times as many occasions at the four exposure durations. All but the 30 iteration value are significantly more frequent than would be expected by chance ($\chi^2 = 16.9$ ($P < 0.001$); $\chi^2 = 0.3$ ($P > 0.05$); $\chi^2 = 7.9$ ($P < 0.01$); $\chi^2 = 10.1$ ($P < 0.01$), respectively). Thus whether Visual errors occur significantly more often than would be expected by chance depends on exposure duration.

Mixed errors also occur more frequently than would be expected by chance. By chance they would be expected to occur on 0.12 (6/52) as many occasions as Unrelated errors: in fact they occurred on 0.48, 0.9, 1.8 and 5 times as often at the four exposure durations ($\chi^2 \geq 120$; $P < 0.001$ in each case).

4.3. Conclusion

The critical finding is that, whatever the stimulus duration, there is a higher rate of Semantic and Visual errors than would be expected by chance. This mimics the finding in the experiments with normal subjects. The pattern of errors resulting from presentation of brief masked stimuli to the undamaged network is qualitatively similar to that which occurs when the network is lesioned and presented with unmasked stimuli. In both cases the relative proportions of Visual and Semantic errors varies between conditions, stimulus duration in one case, lesion site or magnitude in the other. Lesions close to the graphemic units produce more Visual errors; lesions close to the sememe units produce more Semantic errors (Plaut & Shallice, 1993a). In the simulation reported in this paper the relative proportion depends on the stimulus duration before the mask. After 15 iterations Visual and Semantic errors are produced in roughly equal proportions; after 70 iterations Semantic errors outnumber Visual 7:1.

A second finding is that when a stimulus is input briefly to the network and followed by a pattern mask before it has produced a response, failure to make the correct response leads to both omissions and errors, and that the rate of both increases as overall performance drops from 90% correct to 13% correct. This mimics the performance of normal subjects over a range of stimulus durations giving correct performance ranging from 60–20%.

5. General discussion

Deep dyslexic patients, who are assumed to read by the semantic route in isolation, make both semantic and visual errors when reading aloud. Normal subjects also make both semantic and visual errors when trying to identify the meaning of words under degraded input conditions which were intended to make them rely primarily on the semantic route. A connectionist model which maps the spelling patterns of words onto their meaning produces this pattern of errors both when it is lesioned (simulating the patients) and when it is presented with degraded input (simulating the normal subjects). This suggests that the cognitive system contains a mechanism which operates on a related principle to that of the model.

It should be acknowledged that the model incorporates a number of simplifications and approximations relative to the situation for human subjects. The model contains no implementation of early perceptual processes (prior to orthographic input) or later articulatory processes (subsequent to phonological output). It also does not contain the memory, attentional and control mechanisms required to carry out the task of reporting members of an ad hoc semantic category under RSVP conditions. Our approximation of RSVP by naming under brief masked exposure thus does not incorporate errors due to failures of these processes. It also does not permit influences from preceding words in processing the current word (although see Becker, Moscovitch, Behrmann and Joordens (1997); Masson (1995); Plaut and Booth (1999); Plaut and Shallice (1993b) for ways in which such perseverative influences can be instantiated in attractor networks). Finally, the 40-word vocabulary of the model is, of course, vastly reduced relative to that of the subjects. These approximations leave open the possibility that factors outside the implemented model contribute to the error pattern shown by normal subjects and by brain-damaged patients. The fact that, despite these limitations, our model provides a good approximation of both normal and impaired performance suggests that it captures the critical principles underlying human performance in these contexts.

5.1. Attractor basins

The key to understanding why the model produces the ‘signature’ error pattern (the co-occurrence of visual and semantic errors) both when it is lesioned and when it is given inadequate input to produce a correct response lies in the nature of the attractor basins which are created as it learns to map orthography to semantics. In a feed-forward connectionist network (that is, one without recurrent connections) the output is determined by a single computational pass, triggered by the input pattern, through the weighted connections. In contrast, in a network with recurrent connections, activity continues to pass between units for several processing cycles after activation has been initiated by processing of the external input. Thus the final state of the network will not necessarily be that produced by the initial processing of the external stimulus. However, only certain states (i.e. patterns of activity across the units) are stable in recurrent networks. Whatever state arises initially in the network following an input, the interactions between units as activity cycles around the recurrent connections will cause the system to move towards one of the stable patterns. Consequently, these are known as attractors. If the learning phase has been successful the state corresponding to the semantics of each word in the input set will have its own attractor. The set of states which eventually settle to a given attractor pattern is known as its *basin of attraction*. If the state of the network ever enters a basin, further processing (changes produced by activity circulating in the recurrent connections) will take the state of the network towards that represented by the attractor.

5.2. Mapping orthography to semantics

As discussed in the introduction, connectionist networks have a fundamental bias to give similar responses to similar inputs. This is the basis of their ability to generalise, that is, to produce plausible responses to novel inputs. But the similarity bias is problematic in tasks like mapping orthography to semantics where the relationship between input and output is arbitrary. The initial response of a network to visually similar words will be to generate similar semantic patterns. As visually similar words are not usually related in meaning, the similar initial states generated by visually similar words must be separated into the dissimilar patterns corresponding to the dissimilar semantics of the words. Developing attractor basins within semantic space allows the network to do this. Changes caused by subsequent processing with the recurrent connections can cause initially similar patterns to become distinct as the state of the system moves within a basin to its attractor.

Fig. 2 illustrates the behavior of the model as orthography is mapped onto semantics in terms of trajectories through semantic state space. At any instant, the current state of the system (the pattern of activity in the semantic units) is represented as a point in this space (with a position shown in the horizontal and depth dimensions). The height of the point represents the *energy* of the state – the degree to which it violates the constraints imposed by the input and the network's knowledge. (Energy is the inverse of how good the state is as an interpretation of the input). The energy values for every possible pattern of activity form a surface in state space. As activation passes around the recurrent connections and the network settles in response to an input, the point corresponding to the current pattern of semantic activation moves downhill on this surface. This is represented by the paths of the solid lines.

The effect of the orthographic input takes time to diffuse into the network and influence semantics, so the trajectories all start along a similar path. Once orthographic influence starts to have an effect, the trajectories following the input of visually similar words like DOG or LOG remain similar and unlike that following the input of a visually dissimilar word such as CAT. The network must learn to overcome this effect of visual

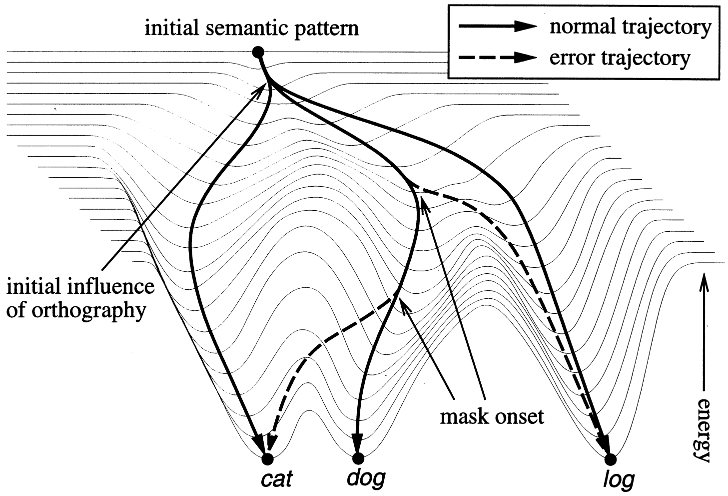

Fig. 2. Hypothetical trajectories in the semantic state space of the system resulting from input of the words DOG, LOG, and CAT. The solid lines show the trajectories following normal, unmasked presentation of the word. The dashed lines show how the trajectory following DOG might be affected if the stimulus was replaced by a mask at two different times after stimulus onset.

similarity in order to derive the correct semantics for the word. The trajectory following the input DOG must end up close to that following the input CAT and far from that following the input LOG. This is accomplished by adjusting the connection weights – and thus the basins in the energy landscape – so that the initially similar trajectories for visually similar but semantically unrelated words are pulled apart and eventually arrive at distant attractors. The attractor basins for any given word must at some point approach the basins of visually related words and at others those of semantically related words. Thus, for example, some parts of the *DOG* basin must be close to the *LOG* basin while other parts must be close to the *CAT* basin. This is the key feature of the model which ensures that it will produce a combination of visual and semantic errors following either input masking or lesioning.

5.3. The co-occurrence of visual and semantic errors

An orthographic mask drives semantic activity in a direction unrelated to that produced by the stimulus word. As a result, the normal trajectory in semantic space begins to be deflected from the point reached when the mask comes on. The crucial point is that the effect of the mask is to *deflect* the trajectory in state-space, not to cause it to jump to a new, random position.[9] The deflection, shown by the dashed

[9] Fig. 2 is somewhat misleading as it shows the mask deflecting the trajectory of the network state over a fixed energy landscape. In fact, the network state always follows the direction of steepest descent; the mask deflects the trajectory by altering the landscape itself, thereby changing the direction of steepest descent. This change is difficult to convey in a static figure so we have opted for the format in Fig. 2.

lines in Fig. 2, may be sufficient to take the trajectory out of the basin which would take it to the correct response and into the surrounding regions of state space. If it enters a basin for a word it will produce an error. Since basins for related words occupy adjacent areas of state space, the basins close to another will be those of words which are related either visually or semantically to the input. Consequently, masking the input produces visual and semantic errors. As far as the network is concerned, words only vary in terms of visual and semantic similarity, so proximity of attractor basins reflects only these dimensions. Presumably in the human lexicon there are other dimensions of similarity. But visual and semantic similarity are sufficiently important that their influence can be seen in the errors which occur when performance breaks down.

Lesioning the network changes the shape of the attractor basins rather than distorting the trajectories. Changing the shape of the *DOG* basin causes it to lose some trajectories to nearby basins, and to capture other trajectories from nearby basins. Since nearby basins are predominantly for words which are related either visually or semantically to the word for which this is the basin, the altered trajectories will mostly go to or come from basins which are either visually or semantically related to the basin they were in originally. As a result, lesioning the network leads to the occurrence of both visual and semantic errors. Thus, although the mechanisms by which lesioning and input masking have their effects are different, they both lead to trajectories being lost to nearby basins, and thus lead to similar error patterns.

5.4. Attractor dynamics and cognitive processes

That the pattern of semantic and visual errors in patients and normal subjects can be accounted for by the same model suggests that it is based on principles which are at work in the human cognitive system. The key principle in the model is that the mapping between successive layers of representation (in this case between representations of orthography and semantics) is achieved through the operation of attractors. In an attractor network, once processing has been started by stimulus input, further processing will continue irrespective of the presence of the stimulus. Were the stimulus present it would, of course, continue to influence the processing, ensuring that the correct response was produced. But in its absence processing does not stop. Whether errors or omissions are produced when the stimulus is removed depends on the nature of the state space. Once a trajectory leaves the basin which it entered under the influence of initial processing of the stimulus its final destiny will depend on what occupies the regions of space around the attractor basins – basins for other words, or basins for spurious attractors which would not produce a response. The first will produce errors, the second omissions.

The model predicts both the increase in errors as performance declines and the kind of errors which are produced. However, there is one way in which its behaviour does not resemble that of human subjects. The production of both errors and omissions in roughly equal proportions continues in the simulation down to stimulus durations which give effectively 0% correct performance. It is unnecessary to do an experiment with normal subjects to demonstrate that this would not be true for them. At very short stimulus durations people would decline to produce any response, so they would produce only omissions, not errors. One reason the model fails to mimic every aspect of human performance in the RSVP task is that the model was not designed specifically to perform this task; it was an existing model taken 'off the shelf', designed for a different task, and tested in a novel context, unanticipated by its original designers. The fact that, in general, it copes so well, is one of its strengths. However, it would be straightforward to adapt the model so that it could produce this aspect of human performance by training the network to develop an attractor around the initial semantic state. To escape from this attractor would require the energy landscape to be altered for a sufficient length of time by a stimulus, implementing a temporal threshold before responding occurred. (Such an attractor would also be useful in ensuring that the network started its processing from more or less the same state each time thus minimizing variations in the final state caused by variations in the initial state).

Interestingly, one other example of the use of a connectionist model to account for similarities in the behaviour of normal and brain-damaged subjects – the work of Mozer and Behrmann (1990) (Behrmann et al., 1991) on neglect dyslexia – also relies on attractor dynamics. In those studies the relevant phenomena concern lexical influences on the performance of patients who also show evidence of a peripheral impairment (to an early attentional mechanism). For example, patients are less likely to exhibit neglect for a word than for a non-word (Sieroff, Pollatsek & Posner, 1988) or for the leftmost of two words if they form a compound word (e.g. COW BOY) than if they do not (e.g. SUN FLY; Behrmann, Moxcovitch, Black & Mozer, 1990). In MORSEL, the model used by Mozer and Behrmann (1990) to account for these effects, bidirectional interactions between higher order orthographic units and lexical/semantic units form attractors for familiar words. When bottom-up perceptual information is degraded by an attentional impairment, these attractors serve to clean up and complete the information, but are more effective in doing so for words (including compounds) than for non-words. Thus the findings from the modelling work on neglect dyslexia concur with those from the current work in suggesting that attractor dynamics play an important role in lexical processing (see also Plaut et al., 1996). The convergence demonstrated in the current work between the behaviour of neuropsychological patients, normal subjects, and a connectionist attractor network provides additional support for the central role of attractor dynamics in cognitive processing.

Why should the system use attractors? Take a feed forward connectionist system which does not have an attractor structure. If a small change occurs to the stimulus, or to any general modulatory input to the weights related, for instance, to the operation of some biochemical pathway, then there will be a corresponding small change in the output. If, however, there is an attractor structure these small changes in the input variables or the weights will lead to correspondingly smaller changes in the output unless the boundary of an attractor basin is crossed. This leads to a relative equivalence in the effect of the output on subsequent systems for a given set of inputs. The system shows an effect analogous to categorical perception.

Moreover this discreteness of effect on other systems gives this type of connectionist system one property in common with symbol systems.

In fact, the relevance of attractors within cognitive psychology goes far beyond lexical processing. They allow best fit matching of input to possible categories to be based on multiple effective dimensions of relevant input. Purely symbolic systems force the theorist to make an unpalatable choice between allowing decomposition within the representations, which produces paradoxes, or rejecting it which seems counter intuitive (see Roelofs, 1997). Attractors allow decomposition into elements – using the distributed representations at one level – but at the same time produce equivalence of effects, within categories, at the next level and also qualitative differences in effects from one category to another at the next. Thus they allow decomposition of semantic representations into elements, but also in their effects on later representations they have some of the characteristics of symbolic representations (see van Gelder (1990) for discussion).

Acknowledgements

We are greatly indebted to Ian Nimmo Smith for statistical advice, to Jackie Chard, the late Kristen Lewis and Frances Watson who ran the experiments, and to the helpful comments of two anonymous referees.

Appendix A

Errors other than those classified as visual in Experiment 1 (Section 3.1). Targets (two on some trials, indicated by + , one on others) are italicised. Responses are in quotes. .... before a response indicates that the subject did not report a single word as the target but suggested some attribute of the target or the general semantic region in which it fell.

- *street* + *spade* ‘something garden-wise like wheelbarrow’
- *pier* ‘....beach or seaside’
- *tiger* ‘hawk’
- *marsh* ‘.... garden’
- *harbour* + *fountain* ‘.... seaside..lighthouse’
- *cloud* ‘helicopter’
- *chapel* + *lawnmower* ‘.... church’
- *spade* + *wagon* ‘shovel’
- *snow* ‘.... storms, weather’
- *deer* ‘stag’
- *cinema* + *chapel* ‘.... church, synagogue’
- *sail* + *lion* ‘yacht’
- *sand* ‘street’
- *river* ‘trailer’
- *fence* ‘animal’

- *pavement* 'house'
- *lorry* + *airport* '.... aeroplanes and runways'
- *rocket* + *pier* '.... sea'
- *buttercup* '.... seaside'
- *submarine* '.... space capsule'
- *cherrytree* 'honeysuckle'
- *submarine* + *waterfall* '.... aeroplanes'
- *street* + *ship* 'yacht or boat'
- *theatre* 'museum sort of building'
- *grass* 'rock'
- *lane* '.... street'
- *rock* 'cobbles'
- *tortoise* + *brook* 'like a ditch'
- *tower* 'tall monument'
- *lane* + *thrush* 'tree of some sort'
- *statue* 'archway'
- *raft* + *satellite* 'astronaut'
- *mist* + *turret* '.... tall'
- *grass* + *cemetery* 'games'

References

Allport, A. (1977). On knowing the meaning of words we are unable to report: the effects of visual masking. In S. Dornic, (pp. 505–534). *Attention and performance. vol. 6*. London: Academic Press.

Becker, S., Moscovitch, M., Behrmann, M., & Joordens, S. (1997). Longterm semantic priming: a computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 1059–1082.

Behrmann, M., Moscovitch, M., Black, S. E., & Mozer, M. C. (1990). Perceptual and conceptual factors in neglect dyslexia: two contrasting case studies. *Brain*, *113*, 1163–1183.

Behrmann, M., Moscovitch, M., & Mozer, M. (1991). Direct attention to words and nonwords in normal subjects and in a computational model: implications for neglect dyslexia. *Cognitive Neuropsychology*, *7*, 213–248.

Caramazza, A., & Hillis, A. (1990). Where do semantic errors come from? *Cortex*, *26*, 95–122.

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

Coltheart, M., Patterson, K., & Marshall, J. (1987). Deep dyslexia since 1980. In M. Coltheart & K. Patterson, & J. Marshall, *Deep dyslexia*, (pp. 407–451). London: Routledge, Kegan Paul.

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.

Friedman, R., & Perlman, M. (1982). On the underlying causes of semantic paralexias in a patient with deep dyslexia. *Neuropsychologia*, *20*, 559–568.

Hillis, A., Rapp, B., Romani, C., & Caramazza, A. (1990). Selective impairments of semantics in lexical processing. *Cognitive Neuropsychology*, *7*, 191–243.

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

Marshall, J., & Newcombe, F. (1966). Syntactic and semantic errors in paralexia. *Neuropsychologia*, *4*, 169–173.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 323.

Morton, J. (1969). The interaction of information in word recognition. *Psychological Review*, *76*, 165–178.

Mozer, M., & Behrmann, M. (1990). On the interaction of selective attention and lexical knowledge: a connectionist account of deep dyslexia. *Journal of Cognitive Neuroscience*, *2*, 96–123.

Mozer, M. C. (1991). *The perception of multiple objects: a connectionist approach*, Cambridge, MA: MIT Press.

Newton, P., & Barry, C. (1997). Concreteness effects in word production but not word comprehension in deep dyslexia. *Cognitive Neuropsychology*, *14*, 481–509.

Plaut, D.C., Booth, J.R. (1999). Individual and developmental differences in semantic priming: empirical and computational support for a single mechanism account of lexical processing. submitted for publication.

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

Plaut, D. C., & Shallice, T. (1993b). Perseverative and semantic influences on visual object naming errors in optic aphasia: a connectionist account. *Journal of Cognitive Neuroscience*, *5*, 89–117.

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*, 56–115.

Potter, M., & Lombardi, L. (1990). Regeneration in the short-term recall of sentences. *Journal of Memory and Language*, *29*, 633–654.

Roelofs, A. (1997). A case for non-decomposition in conceptually driven word retrieval. *Journal of Psycholinguistic Research*, *26*, 33–67.

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

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

Sieroff, E., Pollatsek, A., & Posner, M. I. (1988). Recognition of visual letter strings following injury to the posterior visual spatial attention system. *Cognitive Neuropsychology*, *5*, 427–449.

van Gelder, T. (1990). Compositionality: a connectionist variation on a classical theme. *Cognitive Science*, *14*, 355–384.

Van Orden, G. C. (1987). A ROWS is a ROSE: spelling, sound and reading. *Memory and Cognition*, *15*, 181–198.

Williams, P., & Parkin, A. (1980). On knowing the meaning of words we are unable to report: a confirmation of a guessing explanation. *Quarterly Journal of Experimental Psychology*, *32*, 101–107.