

A Neural Model of Deep Dyslexia

Michael L. Rossen
Department of Psychology
Brown University
Bitnet: Rossen@brownco

Abstract

This paper presents a simulation of the selective deficits and the partial breakdown patterns characteristic of the oral reading performance of deep dyslexics. The most striking symptom of deep dyslexia -- usually considered its defining characteristic -- is the occurrence in oral reading tasks of semantic paralexias: the vocalization of a word semantically related to an isolated, printed target word. The pattern of simulated paralexia errors by the neural model is strongly controlled by the similarity structure of the training set stimuli and, to a lesser extent, the frequency of presentation of stimuli during learning by the model. This result fits well with effects of stimulus type on patterns of paralexia error among deep dyslexics. Further, the model very naturally reproduces the patterns of partial breakdown observed in deep dyslexics, including a slow response time (RT) and within subject variation of response to a particular target word in successive test sessions.

Keywords: neural models, deep dyslexia, similarity structure.

Symptoms and characteristics of deep dyslexia

Loss of a selected subset of cognitive skills often accompanies brain damage. **Deep dyslexia** is an example of a selective symptom-complex resulting from brain damage in which deficits in certain oral reading skills tend to co-occur. The most striking symptom of deep dyslexia -- usually considered its defining characteristic -- is the occurrence in oral reading tasks of **semantic paralexias**: the vocalization of a word semantically related to an isolated, printed target word (e.g., visual target: BEAR; oral response: LION). Deep dyslexics also make visual paralexias, errors where the oral response word graphemically resembles the target word (e.g., visual target: BEAR; oral response: PEAR).

The type of stimulus word used in oral reading tasks strongly affects the likelihood of a paralexia response by a deep dyslexic. Shallice and Warrington (1975) show that nouns of low usage frequency or of low concreteness have a relatively high probability of eliciting a paralexia response, with concreteness having a stronger effect. Syntactic class also has an effect: correct responses are most likely for nouns, with adjectives, verbs and function words (e.g., "is," "to," "and") causing successively higher rates of paralexia error.

The **selectivity** of deep dyslexia is evident in the stimulus effects just mentioned and from the observation that deep dyslexics can often perform tasks analogous to oral reading, but involving different input and/or output modalities, with little or no evidence of degradation in performance. For example, some deep dyslexics can perform almost perfectly in a picture naming task in which the picture stimuli correspond to word stimuli with which the patients make frequent paralexias during oral reading. (Patterson & Marcel, 1977).

Deep dyslexia involves **patterns of partial breakdown**, patterns of a statistical rather than a deterministic nature. In particular:

- (1) Response time (RT) of deep dyslexics in oral reading tasks is slower than for the normal population. In addition, deep dyslexics often make "omissions," in oral reading tasks; that is, they fail to offer any response at all to a printed target stimulus.
- (2) Intra-subject response variation. Consider an oral reading task in which the printed target

stimulus "BEAR" is one of several stimuli presented in random order during a single test session. In successive test sessions, a deep dyslexic might correctly read the word "BEAR" aloud, commit a semantic paralexia ("LION"), and a visual paralexia ("PEAR").

(3) Inter-subject response variation. Two deep dyslexics can make different responses to a particular target word, yet have similar statistical likelihoods of semantic and visual paralexias.

Anderson (1983), Kawamoto (1985), and McClelland & Rumelhart (1986) have all employed neural models to examine various characteristics of brain damaged patients. Kawamoto specifically modeled aspects of deep dyslexia. The simulations in this paper concentrate on some characteristics of deep dyslexia not addressed by these previous studies. Simulations one through three show that a fully trained neural model can simulate paralexic responses and reproduce the patterns of partial breakdown just described. Simulations three and four suggest that the stimulus effects outlined above that underscore the selective nature of deep dyslexia can be interpreted in terms of the similarity structure of word representation (see below) and the frequency with which a word is learned by the model.

The Brain-state-in-a-box (BSB) neural network model

The BSB model (Anderson, 1983) employs a training algorithm (1a) and a classification algorithm (1b). The algorithms employ a *vector of idealized neural activities* to represent *information flowing through* the system. Each vector element can take on a continuous range of values between -1 and 1, representing the minimum and maximum activity levels of a neuron. A matrix of idealized *synaptic weights* inter-connecting the neurons represents *information stored within* the system.

$$f^* (\tau) = A f(\tau) \qquad \text{(basic computation)}$$

A : matrix of synaptic weights
 f : vector of neural activities
 τ : scalar discrete time index

$$A (\tau + 1) = A (\tau) + \gamma [f(\tau) - f^* (\tau)] f^T \qquad \text{(W-H training algorithm) \qquad (1a)}$$

f^T : transpose of f
 γ : scalar learning parameter

$$f (\tau + 1) = \sigma [\alpha f^* (\tau) + \beta f(\tau)] \qquad \text{(BSB classification algorithm) \qquad (1b)}$$

α : scalar feedback parameter
 β : scalar decay parameter
 σ : function that limits activities to region [-1,1]

Training procedure. Training consists of modification of a 50 percent connected weight matrix "A" by repeated application of the Widrow-Hoff (W-H) training algorithm (1a) with each of the members of the training stimulus set (see below). Each stimulus is learned an equal number of times, except in simulation three when learning frequency effects are explicitly examined. To eliminate potential recency effects of training, the order of stimulus presentation is randomized and the learning parameter is tapered as training nears completion (Anderson, 1983). To ensure the robustness of training, noisy versions of the training set vectors are used to test the system. The noise is calculated so that, when added to a training vector to create a test stimulus vector, the test stimulus is always within a cone around the training stimulus whose axial angle is half

Neural Model of Deep Dyslexia

the angle between the training vector and its nearest neighbor in the training set. Training is terminated when the system classifies four sets of noisy training set vectors without error.

The stimuli. A word stimulus is represented as an activity vector that specifies the system's initial state. The classification algorithm (1b) iterates the activity vector through the weight matrix until the vector reaches a corner in activity space, or until until a maximum of 96 iterations is reached. This final activity vector corresponds to a classification of the stimulus vector. In all of the simulations considered here, a 192-dimensional system is used to learn a training set of eight normalized, demeaned stimuli, each pointing to a different corner in activity space. The cosine between two activity vectors defines their *similarity*. Table 1 shows the similarity structure of the training set. Each stimulus vector has a .75 cosine with three *nearest neighbor* stimuli, a .5 cosine with one stimulus and a .375 cosine with the remaining thee stimuli. This idealized similarity structure allows for quantitative analysis of simulated paralexias (see below) as a function of word stimulus similarity.

Performance measures

RT is simulated by the iteration count in the BSB classification algorithm (1b). A final classification by the system is considered an *omission* error if less than 95 percent of the neurons are at their maximum or minimum activity level. If a classification is not an omission, then the system response is defined as the training set vector to which the classification is most similar. If that training set vector was the input stimulus, then the response is considered correct. Otherwise, the response constitutes an erroneous classification, a *simulated paralexia*.

Simulating brain damaged subjects and multiple test sessions

Damage consists of ablation of a randomly chosen, fixed percentage of the synapses of the weight matrix after termination of learning. Different *computer subjects* are simulated using the same trained network with different random ablations. A particular deep dyslexic subject also varies as to which words he/she makes errors on from one test session to the next. To allow this pattern to emerge from the formally deterministic system employed in these simulations, noisy versions of the training set stimuli are employed, as described in the section on training procedure.

Table 1. Similarity structure of training set stimuli.

Training stimulus	Training stimulus							
	1	2	3	4	5	6	7	8
1	1.	.75	.75	.75	.5	.375	.375	.375
2	.75	1.	.75	.75	.375	.5	.375	.375
3	.75	.75	1.	.75	.375	.375	.5	.375
4	.75	.75	.75	1.	.375	.375	.375	.5
5	.5	.375	.375	.375	1.	.75	.75	.75
6	.375	.5	.375	.375	.75	1.	.75	.75
7	.375	.375	.5	.375	.75	.75	1.	.75
8	.375	.375	.375	.5	.75	.75	.75	1.

Simulation one: (When) are errors made by a neural model?

Procedure. Ten Computer subjects are derived from the trained system, with damage ranging from five percent to 95 percent synapse ablation. Each computer subject participates in one test session of classifying noisy versions of the training set stimuli.

Results. Ablation of synapses have observable effects on system dynamics as measured by both RT and error occurrence. Figure 1a shows RT, averaged across the test stimuli, plotted against proportion of synapse ablation. RT reaches the system maximum for at least some of the test stimuli after only a 25 percent ablation. Simulated paralexias occur for only a narrow range of ablation levels (Figure 1b). All responses are correct with 45 percent ablation or less, while only omission errors occur with 75 percent or more synapse ablation.

Discussion. The response time data corresponds well with the tendency of deep dyslexics to respond more slowly than normals. The location of the peak for paralexia errors at 65 percent ablation level underscores the robust nature of distributed memory storage. Increasing maximum iteration number does increase the number of paralexias somewhat, but the basic shape of the curve in Figure 1b remains the same. Simulated paralexias have not been observed with the present stimulus set with less than 55 percent ablation. Pilot studies show that higher or lower cosines among nearest neighbors leads to more or less simulated paralexias, respectively, given a fixed percentage of ablated synapses. No simulated paralexias could be observed at any damage level, given present operational definitions, for stimulus sets in which nearest neighbors had a cosine of only .25. Also, additional training decreases the number of errors for a given amount of damage, but the effect was negligible.

Simulation two: Within subject variability of error

Procedure. This simulation tests the ability of a computer subject with simulated brain damage to give variable classifications to the same test stimulus in successive test sessions. One




Figure 1a




Figure 1b

Neural Model of Deep Dyslexia

Table 2. Simulation 2: Simulated paralexias by one computer subject in 24 test sessions.

Input stimulus	Output corner								paralexias
	1	2	3	4	5	6	7	8	
4	0	15	1	(3)	0	0	0	0	16
5	0	0	0	0	(12)	0	12	0	12

computer subject with 65 percent synapse ablation performs 24 test sessions of identifying noisy versions of the training set stimuli.

Results. Table 2 shows the number and type of simulated paralexias for each test stimulus (omission errors are not included). No errors occur for 6 of the stimuli. 16 simulated paralexias are made on stimulus four; 15 of them converge to stimulus two and one of them to stimulus three. 12 simulated paralexias are made with stimulus five, all to stimulus seven.

Discussion. Simulated paralexia error patterns are variable over the sessions, conforming to the error patterns reported in the literature for deep dyslexic patients. Of the training set members with which simulated paralexias are made, neither of them always produced errors with the system, and one of them produced two different types of paralexias.

Simulation three: Error patterns across computer subjects

Procedure. This simulation examines the effect of similarity between activity vectors on the pattern of simulated paralexias by computer subjects. Because of the the symmetric similarity structure of the training set, paralexia errors should be equally likely for each stimulus. On the other hand, if similarity structure is important, then paralexia errors should be more likely to be classified as near neighbor stimuli. 25 computer subjects each with 65 percent synapse ablation perform one test session of identifying the training set stimuli with no noise added.

Results. Table 3 shows the distribution of simulated paralexias from and to each of the training set stimuli. The most prominent characteristic of these results is that errors only correspond to nearest neighbor stimuli. Within sets of nearest neighbors (stimuli one through four and stimuli five through eight, as can be seen by consulting Table 1), distribution of errors from and to each of the stimuli is almost flat, with the exception of stimulus one to which more errors tended and from which less errors were made.

Discussion. The absence of simulated paralexias to anything but a nearest neighbor stimulus can be interpreted as a tendency for the system to make semantic (as opposed to random) errors when damage is in the form of random synaptic ablation. This prediction assumes that semantically related words are represented in memory as highly correlated patterns of activation. Support for this assumption can be found in the work of Anderson (1983) and Kawamoto (1985). One might speculate further that nouns and adjectives (as opposed to verbs), and words with concrete meanings are represented by activity patterns relatively isolated from the activity patterns of nearest neighbor words. This would serve as an explanation for the syntactic class and concreteness effects found in deep dyslexic oral reading performance.

Table 3. Simulation 3: Classification of test stimuli by 25 computer subjects.

Input stimulus	Output corner								paralexias
	1	2	3	4	5	6	7	8	
1	20	2	1	2	0	0	0	0	5
2	7	13	2	3	0	0	0	0	12
3	3	2	17	3	0	0	0	0	8
4	3	3	2	17	0	0	0	0	8
5	0	0	0	0	15	1	5	4	10
6	0	0	0	0	5	16	2	2	9
7	0	0	0	0	4	2	15	4	10
8	0	0	0	0	2	5	3	15	10
Total paralexias	13	7	5	8	11	8	10	10	72

Simulation four: The effect of variable frequency of stimulus presentation during training.

Procedure. This simulation is concerned with the effects of variable frequency of stimulus presentation during learning on the pattern of simulated paralexias by computer subjects. Toward this end, stimulus one is learned by the system three times as frequently as any of the other stimuli. High frequency of presentation for a particular stimulus causes that stimulus to be learned more quickly than other stimuli. Though this effect is muted somewhat by the error correction nature of the W-H learning algorithm, it is still possible that the high presentation frequency of stimulus one during learning renders it more resistant to error within a damaged system.

A side effect of the skewed presentation frequency is that the system requires relatively more training to learn all the stimuli. As was mentioned in simulation one, one of the results of extra training is a slight increase in stability with respect to damage. Thus, a 75 percent ablation level is used to maximize the number of simulated paralexias. Data is gathered from 100 computer subjects that each perform one test session identifying the training set stimuli with no noise added.

Results. Table 4 shows the distribution of classification errors from and to each of the training set stimuli. A strong tendency exists for simulated paralexias to converge to stimulus one, the stimulus trained with a high frequency of presentation. Conversely, the computer subjects are relatively unlikely to make a simulated paralexia with stimulus one as input.

As in simulation three, no errors at all are made to any but nearest neighbor stimuli. The total number of classification errors made on stimuli five through eight is virtually equal to the number made with stimuli one through four. No stimulus within the set of stimuli five through eight is markedly more or less likely to cause a simulated paralexia.

Neural Model of Deep Dyslexia

Table 4. Simulation 4: Classification of test stimuli by 100 computer subjects. Learning frequency of stimulus one was triple that of the others.

Input stimulus	Output corner								paralexias
	1	2	3	4	5	6	7	8	
1	73	10	8	9	0	0	0	0	27
2	20	49	11	20	0	0	0	0	51
3	34	14	44	8	0	0	0	0	56
4	22	23	7	48	0	0	0	0	52
5	0	0	0	0	60	11	18	11	40
6	0	0	0	0	11	56	15	18	44
7	0	0	0	0	18	16	52	14	48
8	0	0	0	0	18	28	10	44	54
Total paralexias	76	47	26	37	47	55	42	43	373

Discussion. Deep dyslexics tend to make more paralexias with low frequency words; however, this effect is often overshadowed by word concreteness effects. If word concreteness can in part be interpreted as the degree of isolation of a word's activity pattern, then the results of this simulation are in accordance with both of these findings. Moreover, in a pilot study identical to this simulation except that only stimuli one through five were learned by the system, no paralexia errors at all were made with stimulus five, even though stimulus one had a higher learning presentation rate. This model makes the further prediction that when semantic paralexias occur, the paralexia is likely to be of higher frequency in the language than the target word.

References

- Anderson, J. A. (1983). Cognitive and psychological computation with neural models. *IEEE transactions on systems, man, and cybernetics*, SMC-13(5), 799-815.
- Kawamoto, A. H. (1985). Dynamic processes in the (re)solution of lexical ambiguity. Unpublished doctoral dissertation, Providence: Brown University.
- McClelland, J. L., & Rumelhart, D. E. (1986). Amnesia and distributed memory. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol 2. Psychological and Biological Models* (pp. 503-527). Cambridge, MA: Bradford.
- Patterson, K. E., and Marcel, A. J. (1977). Aphasia, dyslexia and phonological coding of written words. *Quarterly Journal of Experimental Psychology*, 29, 307-318.
- Shallice, T., & Warrington, E. K. (1975). Word recognition in a phonemic dyslexic patient. *Quarterly journal of experimental psychology*, 27, 187-199.