

Generalization with Componential Attractors: Word and Nonword Reading in an Attractor Network

David C. Plaut and James L. McClelland

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213-3890
plaut@cmu.edu and jlm@crab.psy.cmu.edu

Abstract

Networks that learn to make familiar activity patterns into stable *attractors* have proven useful in accounting for many aspects of normal and impaired cognition. However, their ability to generalize is questionable, particularly in quasiregular tasks that involve both regularities and exceptions, such as word reading. We trained an attractor network to pronounce virtually all of a large corpus of monosyllabic words, including both regular and exception words. When tested on the lists of pronounceable nonwords used in several empirical studies, its accuracy was closely comparable to that of human subjects. The network generalizes because the attractors it developed for regular words are *componential*—they have substructure that reflects common sublexical correspondences between orthography and phonology. This componentiality is facilitated by the use of orthographic and phonological representations that make explicit the structured relationship between written and spoken words. Furthermore, the componential attractors for regular words coexist with much less componential attractors for exception words. These results demonstrate that attractors can support effective generalization, challenging “dual-route” assumptions that multiple, independent mechanisms are required for quasiregular tasks.

Introduction

Many aspects of language processing can be characterized as *quasiregular*—the relationship between inputs and outputs is systematic but admits many exceptions. An example of such a task is pronouncing English words. Most words are *regular* (e.g., GAVE, MINT) in that they adhere to standard spelling-sound correspondences. In fact, skilled readers can use knowledge of these correspondences to read pronounceable nonwords (e.g., MAVE, BINT). However, they can also correctly pronounce *exception* words (e.g., HAVE, PINT) that violate these correspondences.

*This research is part of an ongoing collaborative project with Mark Seidenberg and Karalyn Patterson, whose contributions we gratefully acknowledge. We also thank Marlene Behrmann for helpful comments on an earlier draft. Financial support for this research is provided by the National Institute of Mental Health (Grants MH00385 and MH47566), the McDonnell-Pew Program in Cognitive Neuroscience (Grant T89-01245-016) and the National Science Foundation (Grant ASC-9109215).

A central question in cognitive science is how best to characterize the language system in order to account for its success at quasiregular tasks. One view (e.g., Pinker, 1991; Pinker & Prince, 1988) is that the language system learns and applies an explicit set of rules, augmented when necessary with a separate enumeration of exceptions. An alternative view, coming out of connectionist or parallel distributed processing research (e.g., Rumelhart & McClelland, 1986), is that the processing of regular and exception items can co-exist within a system that learns to be sensitive to the statistical structure between the inputs and outputs to which it is exposed.

In the specific context of word reading, “dual-route” theorists (e.g., Coltheart, 1985) have claimed that pronouncing exception words requires a “lexical look-up” mechanism that is separate from the grapheme-phoneme correspondence rules that apply to regular words and nonwords. Seidenberg and McClelland (1989, hereafter SM89) have challenged this claim by developing a connectionist network that successfully pronounces both regular and exception words. A major advantage of the connectionist approach is that it provides a more natural account of graded effects of spelling-sound consistency among words (Glushko, 1979) and how this consistency interacts with frequency (Andrews, 1982; Seidenberg, 1985; Seidenberg et al., 1984; Taraban & McClelland, 1987; Waters & Seidenberg, 1985). However, the SM89 network is much worse than skilled readers at pronouncing nonwords (Besner et al., 1990). Since both regular and exception words could be read solely by a lexical procedure, it has been argued that the network’s poor nonword reading is consistent with the dual-route claim that *skilled* reading requires multiple mechanisms (Coltheart et al., in press).

An important connectionist principle lacking in the SM89 simulation (but present in their more general framework for lexical processing) is interactivity. A common way in which interactivity has been employed in networks is in forming *attractors* for particular patterns of activity. In an attractor network, the connection weights cause units to interact in such a way that the initial pattern of activity generated by an input gradually settles to the nearest attractor pattern. If the state of each unit is represented

along a separate dimension in a high-dimensional *state space*, then each attractor pattern corresponds to a particular point within this space, and the set of patterns that settle to it corresponds to a region around it called its *basin of attraction*. In order for the network to perform a task correctly, each input pattern must fall within the attractor basin for the appropriate interpretation.

Interactivity and attractors have played a central role in accounting for a wide variety of phenomena in both normal and impaired word reading (Hinton & Shallice, 1989; McClelland, 1991; McClelland & Rumelhart, 1981; Mozer, 1990; Mozer & Behrmann, 1990; Plaut & Shallice, in press). However, the nature of attractors would appear to be at odds with the form of generalization required for reading nonwords. In an interactive system that has formed attractors for words, one might expect that the input for a nonword would often be captured within the attractor basin for a similar word, resulting in many incorrect responses (lexicalizations). If it is true that attractors cannot support generalization, their applicability in word reading specifically, and cognitive science more generally, would be fundamentally limited.

In this paper, we describe a simulation in which an attractor network learns to pronounce a large corpus of words, including many exception words, and yet also reads pronounceable nonwords as well as skilled readers. The network generalizes because the attractors it develops for regular words are *componential*—they have substructure that reflects common sublexical correspondences between orthography and phonology that also apply to nonwords. This componentiality is facilitated by the use of representations that make explicit the structured relationship between written and spoken words. These findings extend those of Brousse and Smolensky (1989), who found massive generalization in a feedforward autoencoder network trained in a combinatorial environment. In our network, the componential attractors for regular words coexist with much less componential attractors for exception words. The results suggest that, rather than being a hindrance, attractors are a particularly effective style of computation for quasiregular tasks such as word reading.

Simulation

Task definition

The task involves generating the pronunciations of letter strings—that is, mapping orthography to phonology. As a training corpus, we used the 2897 monosyllabic words of SM89, augmented with 101 words missing from that corpus but used as word stimuli in various experiments. This corpus contains almost all monosyllabic words in English. The success of a network in accomplishing this task depends critically on how orthographic and phono-

logical information is represented to the network. To the extent that the representations make explicit the relevant relationships between input and output, the network will learn the task more easily and generalize better.

In phonology, the relevant structure of pronunciations can be described in terms of ordered sets of phonemes. A simple representation would be to have a separate unit for each possible phoneme in each possible position within a pronunciation (e.g., McClelland & Elman, 1986). Unfortunately, this scheme results in poor generalization because the knowledge of when to activate a particular phoneme must be learned separately for each position. At the other extreme, a representation with only a single unit for each phoneme regardless of position would lose information of the relative ordering of phonemes, so that, for instance, /tip/ and /pit/ would be indistinguishable.

However, it turns out that a scheme that involves only a small amount of replication is sufficient to uniquely represent virtually all uninflected monosyllables. By definition, a monosyllable contains only a single vowel, so only one set of vowel units is needed. A monosyllable may contain both an initial and a final consonant cluster, and almost every consonant can occur in either cluster, so separate sets of consonant units are required for the initial and final consonant clusters. The remarkable thing is that this is nearly all that is necessary. The reason is that, within an initial or final consonant cluster, there are strong phonotactic constraints that arise from the structure of the articulatory system. At both ends of the syllable, each phoneme can occur only once, and the order of phonemes is strongly constrained. For example, if the phonemes /s/, /t/ and /r/ all occur in the onset cluster, they must be in that order, /str/. Given this, all that is required to specify a pronunciation is which phonemes are present in each cluster—the phonotactic constraints uniquely determine the order in which these phonemes occur.

There are a small number of cases in which two phonemes can occur in either order within a consonant cluster (e.g., /p/ and /s/ in CLASP, LAPSE). To handle such cases, it is necessary to add units to disambiguate the order (e.g., /ps/). The convention is that, if /s/ and /p/ are both active, they are taken in that order unless the /ps/ unit is active, in which case the order is reversed. To cover the pronunciations in the SM89 corpus, only three such units are required: /ps/, /ts/, and /ks/. Interestingly, these combinations are sometimes treated as single phonemes, called *affricates*, and are sometimes written with single letters (e.g., Greek ψ , English X).

This representational scheme applies almost as well to orthography as it does to phonology because English is an alphabetic language (i.e., parts of the written form of a word correspond to parts of its spoken form). However, the spelling units that correspond to phonemes, called graphemes, are not necessarily single letters (e.g., TH,

Table 1: The orthographic and phonological representations.

| Phonology ^a | |
|------------------------|--|
| onset | s b p d t g k f v z T D S Z l r w m n h y |
| vowel | a @ e i o u A E I O U W Y A |
| coda | r l m n N b g d p s k s t s f v p k t z S Z T D |
| Orthography | |
| onset | Y S P T K Q C B D G F V J Z L M N R W H U C H G H G N G U P H P S Q U R H S H T H T S W H |
| vowel | E I O U A Y A I A U A W A Y E A E E E I E U E W E Y I E O A O E O I O O O U O W O Y U E U I U Y |
| coda | H R L M N B D G C X F V J S Z P T K B B C H C K D D D G F F G G G H G N G U K S L L N G N N P H P P P S Q U R R S H S L S S T C H T H T S T T Z Z E E S E D |

^a/a/ in POT, /@/ in CAT, /e/ in BED, /i/ in HIT, /o/ in DOG, /u/ in GOOD, /A/ in MAKE, /E/ in KEEP, /I/ in BIKE, /O/ in HOPE, /U/ in BOOT, /W/ in NOW, /Y/ in BOY, /A/ in CUP, /N/ in RING, /S/ in SHE, /Z/ in BEIGE, /r/ in THIN, /d/ in THIS.




Figure 1: The architecture of the network. Arrows represent complete connectivity between or within layers.

PH). For this reason, the spelling-sound regularities of English can be captured most effectively if the orthographic units represent graphemes instead of letters. To be consistent with the treatment of affricates in phonology, whenever a multiletter grapheme is present (e.g., TH) we also activate its components (e.g., T and H).

Table 1 presents the details of the phonological and orthographic representations used in the simulation. A total of 108 graphemes and 57 phonemes are required to represent all of the words in the training corpus.

Network architecture

Figure 1 depicts the architecture of the network. In addition to the grapheme and phoneme units, the network contains a single layer of 100 hidden units. These hidden units are connected to each grapheme unit, each phoneme unit, and each other. The phoneme units are also fully connected to each other. All connections are bidirectional, and the weights are initialized to small random values. Including bias connections, the network has a total of 23,203 connections. The units are standard sigmoidal units with real-valued output ranging between 0.0 and 1.0.

Training procedure

The network was trained using a continuous version of back-propagation through time (Pearlmutter, 1989). An input string is presented to the network by clamping the states of the appropriate grapheme units. The network is given a fixed amount of time to process the input, during

which the units change their states gradually in response to their net input from other units. The network's performance was measured by the cross-entropy (Hinton, 1989) of the phoneme units' activity with their desired activity. For the purposes of training, the weighting of this error measure was gradually increased during settling, pressuring the network to be correct as quickly as possible. Also, the network was halted once it succeeded in activating all phonemes to within 0.2 of their correct values. The time to halt was used as a measure of the naming latency of the network.

In addition to the monosyllabic words, the training corpus included training patterns consisting of each single grapheme and the corresponding phoneme, because many children are explicitly taught these correspondences when learning to read. Rather than present each word with a probability proportional to its frequency of occurrence (Kucera & Francis, 1967) and update the weights immediately, we accumulated the error derivatives for the training cases, each weighted by its frequency, before changing the weights. This enabled the learning rates on each connection to be adapted independently during training (Jacobs, 1988; but see Sutton, 1992, for a recently developed online version).

Testing procedure

The ordering of phoneme units, shown in Table 1, embodies the relevant phonotactic constraints on the pronunciation of English monosyllables. Accordingly, the response of the network to any orthographic input can be read off simply by scanning the phonemes in left-to-right order and concatenating all active phonemes (i.e., with activity above 0.5). Because each pronunciation must contain exactly one vowel, only the most active vowel is included in the response. If any affricate unit is active, the order of the corresponding phonemes is reversed.

Results

After 3200 sweeps through the training corpus, the network correctly pronounced all but 10 of the words (99.7%

correct). The missed words were all low-frequency exception words (e.g., BEIGE, SIOUX) and many of the incorrect responses to these were regularizations (e.g., SIEVE \Rightarrow /sev/, SPA \Rightarrow /sp@/). Also, the naming latencies of the network on frequency-matched sets of 48 regular words and 48 exception words (Taraban & McClelland, 1987) exhibit the standard effects of frequency (high 1.82 vs. low 1.93, $F_{1,92}=7.18$, $p<.01$), consistency (regular 1.67 vs. exception 2.09, $F_{1,92}=95.5$, $p<.0005$), and their interaction ($F_{1,92}=3.63$, $p=.06$) observed in human subjects.

However, our main concern is with how well the network generalizes—that is, how well it reads pronounceable nonwords. We tested the network on three lists of nonwords from two experimental studies. The first two lists come from an experiment by Glushko (1979), in which he compared subjects' reading of 43 nonwords derived from regular words (e.g., HEAN from DEAN) with their reading of 43 nonwords derived from exception words (e.g., HEAF from DEAF). The third list of nonwords comes from a study by McCann and Besner (1987), in which they compared performance on a set of pseudohomophones (e.g., BRANE) with a set of 160 control nonwords (e.g., FRANE). We used only their control list in the present investigation because we believe pseudohomophone effects are mediated by aspects of the reading system that are not implemented in our simulation. For the purposes of an initial comparison, we considered the response to a nonword to be correct if it was "regular", as defined by adhering to the grapheme-phoneme correspondence rules as outlined by Venezky (1970).

On the Glushko regular nonwords, both the subjects and the network are near perfect (subjects: 93.8% correct; network: 97.7% correct). Performance on the McCann and Besner control nonwords is somewhat worse (subjects: 88.6% correct; network: 88.8% correct). This is not surprising as the list contains a number of orthographically unusual nonwords (e.g., JINJE, VAWX) that are more difficult to pronounce, both for subjects and for the network.

On the Glushko exception nonwords, both subjects and the network frequently fail to produce the correct (regular) pronunciation, with the network being somewhat worse (subjects: 78.3% correct; network: 69.8% correct). However, to understand this discrepancy, we must reevaluate how a "regular" response is defined. Consider the nonword GROOK. What counts as regular depends on whether we consider the context in which the vowel occurs. OO is most frequently pronounced /u/ as in BOOT, so one possibility is that /grUk/ is the regular pronunciation for GROOK. However, final OOK is almost always pronounced /uk/ as in TOOK, so perhaps /gruk/ should be regular. Thus far, we have considered /grUk/ to be correct and /gruk/ to be an error. However, subjects are sensitive to context in which vowels occur, as evidenced by their much

poorer performance on the Glushko exception nonwords than on the regular nonwords. In fact, Glushko found that 80% of subject's non-regular responses to exception nonwords were consistent with some other pronunciation of the nonword's body that occurs in the Kucera and Francis (1967) corpus, leaving only 4.1% of the responses as actual errors. Similarly, in the network, 84.4% of the non-regular responses to exception nonwords match some other pronunciation in the training corpus for the same body, and over half of these were the most frequent pronunciation of the body. Only 4.7% of the network's responses were actual errors.

It must be acknowledged that the network's behavior does not perfectly match that of subjects. Occasionally, the network makes a frank mistake that subjects would not make, such as omitting a phoneme or pronouncing PH as /p/ instead of /f/. Nonetheless, the network is essentially perfect at pronouncing words, and generalizes to reading nonwords about as well as skilled readers. Thus, these results constitute a direct challenge to dual-route claims that skilled reading requires multiple, independent mechanisms (Coltheart et al., in press).

Network analyses

The network's success at word reading demonstrates that, through training, it has developed attractors for the pronunciations of words. How then is it capable of reading nonwords with novel pronunciations? Why isn't the input for a nonword (e.g., MAVE) captured by the attractor for an orthographically similar word (e.g., GAVE, MOVE, MAKE)? We carried out three analyses of the network to better understand its ability to read nonwords. Because nonword reading involves recombining knowledge derived from word pronunciation, we were primarily concerned with how separate parts of the input contribute to (1) the correctness of parts of the output, and (2 and 3) the hidden representation for the word. The analyses involved stimuli from Taraban and McClelland (1987), all of which the network reads correctly.

The first analysis involved measuring how sensitive the activity in each phonological cluster is to changes in the activity of each orthographic cluster. For each word, the activity of the grapheme units in a particular orthographic cluster were gradually reduced until, when the network was rerun, the phonemes in a particular phonological cluster were no longer correct (i.e., at least one phoneme was on the wrong side of 0.5). This "boundary" activity level measures how important input from a particular orthographic cluster is to the correctness of a particular phonological cluster; a value of 1.0 means that the graphemes in that cluster must be completely active; a value of 0.0 means that the phonemes are completely insensitive to the graphemes in that cluster. The boundary level can also be interpreted as the radius of the word's attractor basin along

Table 2: The sensitivity of each phonological cluster to each orthographic cluster in regular and exception words.

| | Orth. Cluster | Phon. Cluster | | |
|-----------------|---------------|---------------|-------|-------|
| | | onset | vowel | coda |
| Regular Words | onset | 0.491 | 0.000 | 0.000 |
| | vowel | 0.000 | 0.439 | 0.023 |
| | coda | 0.000 | 0.031 | 0.570 |
| Exception Words | onset | 0.520 | 0.433 | 0.019 |
| | vowel | 0.000 | 0.492 | 0.044 |
| | coda | 0.012 | 0.428 | 0.503 |

Note: Words from Taraban and McClelland (1987).

a particular direction in state space.

Table 2 presents the average boundary activity levels for each combination of orthographic and phonological clusters in regular and exception words. First consider the regular words. The diagonal entries are all large, indicating that each phonological cluster is quite sensitive to activity in the corresponding orthographic cluster. In contrast, phonological clusters are almost completely insensitive to activity in the remaining clusters, although there is a slight mutual dependency between the vowel and coda. Thus, an alternative onset (e.g., in a nonword) can be substituted without affecting the pronunciation of the body. Essentially, the attractor basin for a regular word consists of three separate, orthogonal sub-basins: one for the onset, one for the vowel, and one for the coda. When the word is presented, the network settles into the region in state space where these three sub-basins overlap, corresponding to the word's pronunciation. However, each sub-basin can apply independently, so that so-called "spurious" attractor basins for nonwords exist where the sub-basins for parts of words overlap.

This componentiality arises directly out of the degree to which the network's representations make explicit the structure of the task. By minimizing the extent to which information is replicated, the representations condense the regularities between orthography and phonology. Only small portions of the input and output are relevant to a particular regularity, allowing it to operate independently of other regularities.

Returning to Table 2, the attractor basins for exception words are, by contrast, far less componential than those for regular words. In particular, achieving the correct vowel phoneme depends on the entire orthographic input. In fact, most exception words are not regular precisely because they contain an unusual vowel pronunciation. Thus, the network develops noncomponential attractor basins for words when necessary. In this way, the network can pronounce exception words and yet still generalize well to nonwords.

The first analysis establishes the componentiality of the

attractors for regular words behaviorally, but provides little insight into how this componentiality is implemented by the hidden units. One possibility, consistent with dual-route theories, is that the network has partitioned itself into two sub-networks, one that reads regular words, and another that reads exception words. If this were the case, we would expect some hidden units to contribute to exception words but not to nonwords, while others would contribute to nonwords but not to exception words. We measured the contribution a hidden unit makes to pronouncing a letter string by how much the error in pronouncing the string increases when the unit is removed from the network. If the network had partitioned itself, there would be a negative correlation across hidden units between the number of exception words and the number of nonwords to which each hidden unit makes a substantial contribution (greater than 0.025). In fact, there is a moderate *positive* correlation between the numbers of exception words and nonwords to which hidden units contribute ($r=.43$, $p<.0001$). Thus, some units are more important for the overall task and some are less important, but the network has not partitioned itself into one system that learns the rules and another system that learns the exceptions.

The questions remains, then, as to how the network—as a single mechanism—implements componential attractors for regular words (and nonwords) and noncomponential attractors for exception words. The final analysis attempts to characterize the degree to which hidden representations for regular vs. exception words reflect the differences in the componentiality of their attractors. Specifically, we determined to what extent the contribution that a cluster makes to the hidden representation depends on the context in which it occurs—this should be less for words with more componential representations. The contribution of a cluster in a particular context was measured by the difference between two hidden representations: the one generated by the context with the cluster, and a baseline representation generated by the context alone. For each word, we computed the correlation of the contribution that each orthographic cluster makes in the context of the entire word with its contribution when presented in isolation.¹ A high correlation indicates that the contribution of a cluster to the hidden representation is independent of the presence of other clusters, and hence, reflects a high degree of componentiality.

For each consonant cluster, the average correlation for regular words is significantly higher than for exception words [onset: regular mean 0.776 (sd 0.059) vs. exception mean 0.731 (sd 0.072), $F_{1,92}=10.6$, $p<.002$; coda: regular mean 0.732 (sd 0.130) vs. exception mean 0.643 (sd 0.210), $F_{1,94}=6.25$, $p<.015$]. For vowel clusters, there is no significant difference [regular mean 0.703 (sd 0.126)

¹For a cluster in isolation, the context without the cluster is no input at all.

vs. exception mean 0.694 (sd 0.104), $F < 1$). This may be because all words contain a vowel, so that the word presented without a vowel may still activate the appropriate vowel phoneme. Thus, at least with respect to consonant clusters, the representations of regular words are more componential than those of exception words. However, what is surprising is that the average correlations for exception words, though lower than those of regular words, are still quite high, and there is considerable overlap between the distributions. Furthermore, the representations for regular words are not completely componential, given that their correlations are significantly less than 1.0.

Apparently, the hidden representations do not strongly distinguish between regular and exception words. Rather, these representations seem to retain information about the local content of individual clusters, while also capturing some higher-order orthographic structure. The presence of this higher-order structure is what makes the representation of clusters in both regular and exception words somewhat sensitive to context in which they occur. At the phonological layer, information from the hidden representation about individual clusters supports componential attractors for regular words, while the higher-order structure supports less-componential attractors for exceptions.

Conclusion

Interactivity, and its use in implementing attractors, is an important computational principle in connectionist accounts of a wide range of cognitive phenomena. However, the tendency of attractors to capture similar patterns would appear to make them inappropriate for cognitive tasks, such as word reading, which require novel responses to novel inputs. The current research shows, to the contrary, that using representations that condense regularities between inputs and outputs leads to the development of attractors with componential structure that supports effective generalization. At the same time, the network can also learn to develop noncomponential attractors for items that violate the regularities in the task. In this way, attractors provide an effective means of capturing both the regularities and the exceptions in a quasiregular task. Given this demonstration, the claim of dual-route theorists that such tasks require multiple mechanisms appears unfounded.

References

- Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory and Cognition*, 10:565-575.
- Besner, D., Twilley, L., McCann, R. S., & Seergobin, K. (1990). On the connection between connectionism and data: Are a few words necessary? *Psychological Review*, 97(3):432-446.
- Brousse, O. & Smolensky, P. (1989). Virtual memories and massive generalization in connectionist combinatorial learning. In *Proceedings of the 11th Annual Conference of the Cognitive Science Society* (pp. 380-387). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Coltheart, M. (1985). Cognitive neuropsychology and the study of reading. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and Performance XI* (pp. 3-37). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Coltheart, M., Curtis, B., & Atkins, P. (in press). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*.
- 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.
- Hinton, G. E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40:185-234.
- Hinton, G. E. & Shallice, T. (1989). Lesioning a connectionist network: Investigations of acquired dyslexia (Technical Report CRG-TR-89-3). Toronto, Ontario, Canada: University of Toronto, Department of Computer Science.
- Jacobs, R. A. (1988). Increased rates of convergence through learning rate adaptation. *Neural Networks*, 1:295-307.
- Kucera, H. & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- McCann, R. S. & Besner, D. (1987). Reading pseudohomophones: Implications for models of pronunciation and the locus of the word-frequency effects in word naming. *Journal of Experimental Psychology: Human Perception and Performance*, 13:14-24.
- McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology*, 23:1-44.
- McClelland, J. L. & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18:1-86.
- McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part I. An account of basic findings. *Psychological Review*, 88(5):375-407.
- Mozer, M. C. (1990). *The perception of multiple objects: A connectionist approach*. Cambridge, MA: MIT Press.
- Mozer, M. C. & Behrmann, M. (1990). On the interaction of selective attention and lexical knowledge: A connectionist account of neglect dyslexia. *Journal of Cognitive Neuroscience*, 2(2):96-123.
- Pearlmutter, B. A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation*, 1(2):263-269.
- Pinker, S. (1991). Rules of language. *Science*, 253:530-535.
- Pinker, S. & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28:73-193.
- Plaut, D. C. & Shallice, T. (in press). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*.
- Rumelhart, D. E. & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP research group (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models* (pp. 216-271). Cambridge, MA: MIT Press.
- Seidenberg, M. S. (1985). The time course of phonological code activation in two writing systems. *Cognition*, 19:1-10.
- Seidenberg, M. S. & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96:523-568.
- Seidenberg, M. S., Waters, G. S., Barnes, M. A., & Tanenhaus, M. K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning and Verbal Behaviour*, 23:383-404.
- Sutton, R. S. (1992). Adapting bias by gradient descent: An incremental version of Delta-Bar-Delta. In *Proceedings of the 11th National Conference on Artificial Intelligence*.
- Taraban, R. & McClelland, J. L. (1987). Conspiracy effects in word recognition. *Journal of Memory and Language*, 26:608-631.
- Venezky, R. L. (1970). *The structure of English orthography*. The Hague: Mouton.
- Waters, G. S. & Seidenberg, M. S. (1985). Spelling-sound effects in reading: Time course and decision criteria. *Memory and Cognition*, 13:557-572.