Seidenberg, M.S., (2012). Computational models of reading: connectionist and dual-route approaches. In M. Spivey, K. McRae, & M. Joanisse (Eds.), Cambridge Handbook of Psycholinguistics. Cambridge University Press, pp. 186-203.

CHAPTER 9

Computational Models of Reading

Connectionist and Dual-Route Approaches

Mark S. Seidenberg

Word recognition in reading is one of the most extensively studied topics in cognitive psychology-neuropsychology-neuroscience. Computational models are among the tools available for developing a deeper understanding of cognitive phenomena, and they have played a particularly important role in research on normal and disordered reading. My goal in this chapter is to review some of the historical background for the development of present-day computational models of word reading, and then consider two influential contemporary approaches: the Dual Route Cascade (DRC) approach (Coltheart et al., 1993; 2001) and the triangle connectionist approach (Harm and Seidenberg, 1999, 2004; Plaut et al., 1996; Seidenberg and McClelland, 1989). These are not the only computational models of reading in the marketplace, but I have focused on them for several reasons. One is because it isn't possible to review all the existing models in the available space. Another is because these two frameworks represent very different approaches to computational modeling and so the contrast between them is informative. Finally, these models have engaged a somewhat broader range of phenomena than other, more specialized models (e.g., ones that focus on a particular task, such as lexical decision (e.g., Wagenmakers et al., 2004).[1]

1 Background

Contemporary computational models of word reading derive from several progenitors that predate the simulation modeling era. The patriarch of the family is perhaps Morton's (1969) logogen model, which was

1 Most people refer to models of word recognition, but I dislike the term because *recognition* seems inconsistent with the kinds of processes that occur in PDP models in which there is no privileged moment of lexical access. It also seems awkward to use this term in reference to tasks as different as determining a word's meaning, reading it aloud, and spelling it. I prefer the term *models of word reading*. This ignores the fact that people also read nonwords, but seems the better alternative, given the artificiality of nonword reading. I am also aware that the term *DRC* was not introduced until Coltheart et al. (2001); however, the model presented in that article was an extension of the 1993 model and so I have used the term to apply to both.

the first, to my knowledge, to utilize the metaphor of a lexicon consisting of entries (logogens, later termed *localist nodes*) corresponding to individual words. The logogen for a word encoded (responded to) information about its visual, semantic, and acoustic codes. These codes were represented by sets of attributes (or "features"). When a word was processed, the numerical value associated with the logogen that encoded any of the word's visual, semantic, or acoustic attributes was incremented. It would be convenient to term this "activating" a logogen, but this term was not introduced until later. Other logogens that encoded these features would also be incremented (e.g., the /b/ in *book* would increment the values for other words containing this phoneme), leading to what would later be called "competition among partially activated alternatives." The "winning" logogen was determined by Luce's choice rule (Luce, 1959). The resemblances to later localist connectionist models (e.g., McClelland and Rumelhart, 1981; the lexical route in Coltheart et al.'s 2001 DRC model, both discussed later in this chapter) are obvious. It is worth pausing for a moment to recognize the sheer number of essential concepts that were introduced in this pioneering work and incorporated in modified or relabeled form in later work.

A second important ancestor was the model developed by Marshall and Newcombe (1973), which introduced the idea of different word processing routines (visual, phonological, semantic). This general framework later became known as the dual-route model (Baron and Strawson, 1976; Coltheart, 1978). Based on a logical analysis of the properties of English spelling, Marshall and Newcombe deduced that people's ability to read aloud irregularly pronounced words (such as *pint* and *give*) as well as unfamiliar, nonce words (such as *nust*) must involve two mechanisms, later termed the *lexical* and *sublexical* routes, respectively. Whereas Morton provided a mathematical characterization of the major tenets of his theory, dual-route models were stated informally in a quasi-computational information processing language. Again it is worth pausing to acknowledge the extent to which the later computational versions of the dual-route model (i.e., DRC) adhere to these earlier informal proposals.

The dual-route model was the dominant theory of word recognition through the 1980s. As in Marshall and Newcombe's original work, much of the motivation for the model came from studies of patients whose brain injuries impaired their reading in different ways. For example, patients termed "phonological dyslexics" were relatively more impaired at reading nonwords (e.g., *nust*, *faige*) than "regular" words such as *must* or "exception" words such as *have*. This selective impairment of nonwords was said to imply damage to the nonlexical mechanism responsible for generating the pronunciations of novel letter strings, which Coltheart (1978) construed as "grapheme–phoneme correspondence rules" (GPCs). Conversely, patients termed "surface dyslexics" were relatively more impaired in reading exception words than regular words or nonwords. This implied damage to the lexical naming mechanism responsible for word-specific information. This double dissociation was thought to provide strong evidence for the two independent, isolable naming routines that gave the dual-route model its name. Several variants of the basic dual-route scheme were introduced in attempts to capture other types of impaired reading such as deep dyslexia and letter-by-letter reading (see Coltheart et al., 1980; Patterson et al., 1985). The methodology emphasized the importance of relatively clear cases in which one or another component of the model was impaired while others remained largely intact. Such cases might be highly atypical (insofar as such extreme patterns were rarely observed) but nonetheless highly informative. The co-occurrence of deficits was thought to be less informative than selective impairments; a messy brain lesion might affect several components of the reading system, creating co-occurring but otherwise unrelated deficits.

Such was the state of affairs in the modeling of word recognition circa 1987. Model development was closely tied to the methods, goals, and types of data and theorizing associated with cognitive neuropsychology. Informal "box and arrow" information processing models were proposed as theories of the "functional architecture of the cognitive system;" the pattern of preserved and impaired performance on a set of tasks was used to identify the locus of a brain-injured patient's lesion or lesions within this system. The dual-route model of reading aloud provided the paradigmatic example of the approach. Over time, Marshall and Newcombe's relatively simple and elegant model evolved into more complex systems that dealt with a broader range of behavioral phenomena (e.g., agnosia, dysgraphia). Aside from a brief detour into "analogy" models (Kay and Marcel, 1981), however, the models retained the two naming mechanisms of the dual-route approach at their theoretical core.

1.1 *Transitions*

In the late 1980s, several developments conspired to disrupt the status quo. One was dissatisfaction in some quarters about the types of theories that resulted from the functional architecture approach. I summarized many of these concerns in Seidenberg (1988). The box and arrow models were static diagrams that left many details about how information was represented and processed unspecified. The diagrams were not well-suited to characterizing the time course of events that occurred in performing a task such as pronouncing a letter string aloud. The absence of these types of information made it difficult to assess the validity of claims pitched at the level of the functional architecture; it was difficult to determine if a proposed mechanism would produce the intended results. The models were unconstrained in the sense that new structures and operations could be added in response to specific phenomena, even the behavior of a single patient. This lent them an ad hoc character. The lack of specificity and the lack of constraints on the formalism created two problems. One is the familiar lack of falsifiability: If it couldn't be determined if proposed mechanisms would work in intended ways, it couldn't be determined if they would *fail* to work. Moreover, it was too easy to adjust models to fit specific data patterns by adding new modules or processing assumptions. A second, equally important concern is that such models could account for both phenomena that do occur and ones that do not (see Seidenberg, 1993 for discussion). It isn't sufficient that a model account for particular data patterns; given an unconstrained modeling language, and a narrow focus on particular phenomena, any pattern can be mimicked. Considerable explanatory power is gained if a model can simultaneously account for why other outcomes are *not* observed. Whereas falsifiability involves finding data that disconfirm a theory, this second criterion involves showing that a model correctly rules out data that do not occur. The box and arrow models did not meet the latter challenge (Seidenberg, 1988). Rather, they seemed dangerously close to redescriptions of phenomena in a pseudo-computational language.

Other concerns arose about methodology. Numerous battles broke out about best practices. There was an extended debate about the value of single case studies versus analyses of groups of subjects (cf. Caplan, 1988; Caramazza, 1986, and others). Unsurprisingly perhaps, the debate did not yield a consensus, but, utility of group studies aside, one effect was to strengthen arguments in favor of the importance of individual case studies. Eventually, however, additional concerns arose about the interpretation of case studies (Plaut, 1995; Woollams et al., 2006). Another question was whether the methodology was powerful enough to converge on the correct theory of the functional architecture. For example, there was extended debate about whether the functional architecture includes a single semantic system or multiple ones (cf. Caramazza et al., 1990; Shallice, 1993). The conjunction of informal modeling and behavioral data was not sufficient to converge on a clear answer. Concerns were raised about the linkage between data and theory, specifically the kinds of inferences that could be drawn from patient data (Farah, 1994). Many people took issue with the classical interpretation of double dissociations, calling into question whether such dissociations provide airtight evidence for independent components in the functional architecture (cf. Dunn and Kirsner, 2003; Juola and Plunkett, 2000; Plaut, 1995; Van Orden et al., 2001). Such dissociations might arise from other sources: For example, they might represent different points on a distribution of effects created by a combination of factors including different types or degrees of damage to a single system, different degrees of recovery of function or effects of different types of remediation, or different effects of a given type of pathology due to premorbid individual differences. Many of these concerns arose in connection with models of reading, however, they applied more broadly to the functional architecture approach.

1.2 *Enter computational models*

There were several responses to these challenges. Some controversies merely flared and burned out, as researchers lost interest in issues that became bogged down without clear resolution (see, for example, the debate about deficits in access versus representation; Rapp and Caramazza, 1993). Some researchers continued to rely on the traditional approach (see, for example, Miozzo, 2003 and Rapp and Caramazza, 2002, which still use individual case studies and the classical logic governing the identification of components of the functional architecture). However, some researchers took up the challenge of developing models that were more fully specified at the level of how knowledge is represented and processed. The method for achieving this level of mechanistic detail was the implementation of simulation models.

The simulation modeling methodology had been introduced to psychology some years earlier by Newell and Simon (1963). Their models mainly addressed principles of reasoning and problem solving. The pioneering application of this approach to phenomena related to reading was the McClelland and Rumelhart (1981) model. The main purpose of this model was to examine the role of interactive processing in perception; they happened to use letter and word processing as a domain in which to explore this idea. The model demonstrated how interactivity between levels of information (words, letters, features) allowed the system to converge on the identity of a stimulus. In doing so, the model simulated some counterintuitive behavioral phenomena. Letters are easier to identify in the context of words and pseudowords than in isolation (the word superiority effect; Reicher, 1969). Why would the simpler stimulus (a single letter) be harder to identify than when it occurred as part of a more complicated stimulus (a word or pseudoword)? The interactions between top-down and bottom-up flow of activation in the McClelland and Rumelhart model provided the answer.

McClelland and Rumelhart (1981) did not present a general model of reading. There was no phonology or semantics, even though these codes are crucial to reading. The model was limited in other respects: the levels of representation, the nature of the connections between and within levels, the values of the parameters governing the spread of activation, and other properties of the model were hand-wired. Although sufficient for the purpose of exploring the concept of interactivity, these limitations raised questions that later models of word reading attempted to address.

The McClelland and Rumelhart model provided a vivid demonstration of the value of the simulation modeling framework. The ideas were worked out in sufficient detail to be implemented as a computer program. The program simulated detailed aspects of an empirical literature (on word superiority effects) and generated testable predictions. It instantiated general concepts (e.g., interactive activation) that were relevant to reading but not specifically tailored or limited to the target phenomena. The 1981 article is now among the most highly cited in the history of *Psychological Review*. The presentation of a working model was surely a large part of their appeal; most readers probably have not been that interested in word and letter processing. Moreover, Adams (1979) had also recognized the key idea that the word superiority effect might result from feedback from word to letter levels, and applied this idea to behavioral data, but she lacked a computational model.

Given the availability of simulation modeling techniques, and the apparent limitations of the more informal style of modeling, it was not long before researchers began implementing computational models that addressed the phenomena concerning normal and impaired word and nonword naming that had been the focus of so much research within the functional architecture approach. Research following the McClelland and Rumelhart model branched in three directions. One was the implementation of connectionist models based on the parallel distributed processing approach developed by Rumelhart, Hinton, and McClelland (1986). Another was Coltheart and colleagues' Dual Route Cascade implementation of the dual-route model, which incorporated a McClelland and Rumelhart-style interactive activation model as the lexical route. Finally, some researchers continued to use variants of McClelland and Rumelhart's model to study issues such as the processing of orthographic information and particularly the role of orthographic neighborhoods in performing lexical decisions (e.g., Grainger and Jacobs, 1994). As noted earlier, I will focus on the first two lines of research; for a review of the third, see Grainger and Jacobs (1998).

2 The PDP models

In a 1989 article, McClelland and I outlined the general theoretical framework illustrated in Figure 9.1. The framework assumed that words are represented by patterns of activation over units representing spelling, sound, and meaning. Context units were included in recognition of the fact that words typically occur in contexts that affect their meanings (e.g., *the* rose versus *he* rose). The implemented model computed phonological

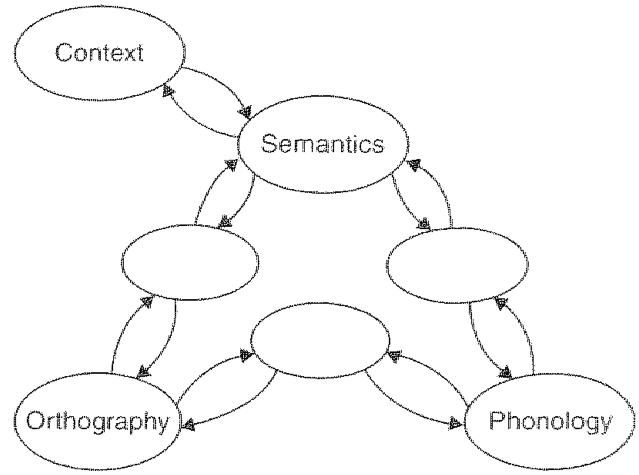

Figure 9.1. General framework for lexical processing introduced by Seidenberg and McClelland (1989). Reprinted with permission.

codes from spellings plus a recreation of the input orthographic pattern. Although the article and implemented model focused on word reading, the framework was intended to represent core processes involved in many uses of words. *Reading* is the process of computing a meaning (or pronunciation) from print. *Spelling* is computing from sound or meaning to print. *Listening*: phonology to meaning. *Production*: meaning to phonology. Seidenberg and McClelland (1989) focused on issues concerning the computation of phonology from print. This seemed like an interesting learning problem insofar as the correspondences between spelling and sound are systematic but include many exceptions, which differ from the central tendencies in differing degrees. I coined the term *quasiregular* to refer to knowledge systems with this character. My interest in the issue arose out of empirical studies of child and adult reading; McClelland's arose out of studies of the past tense in English, which is also quasiregular.

Figure 9.2 illustrates several reading models that have been implemented utilizing principles taken from the PDP framework. The figure brings out the fact that these models bear a family resemblance to each other: They overlap in many respects but no two are identical. What links the models is adherence to PDP principles. None of the models fully incorporate all of the principles, although some principles (e.g.,

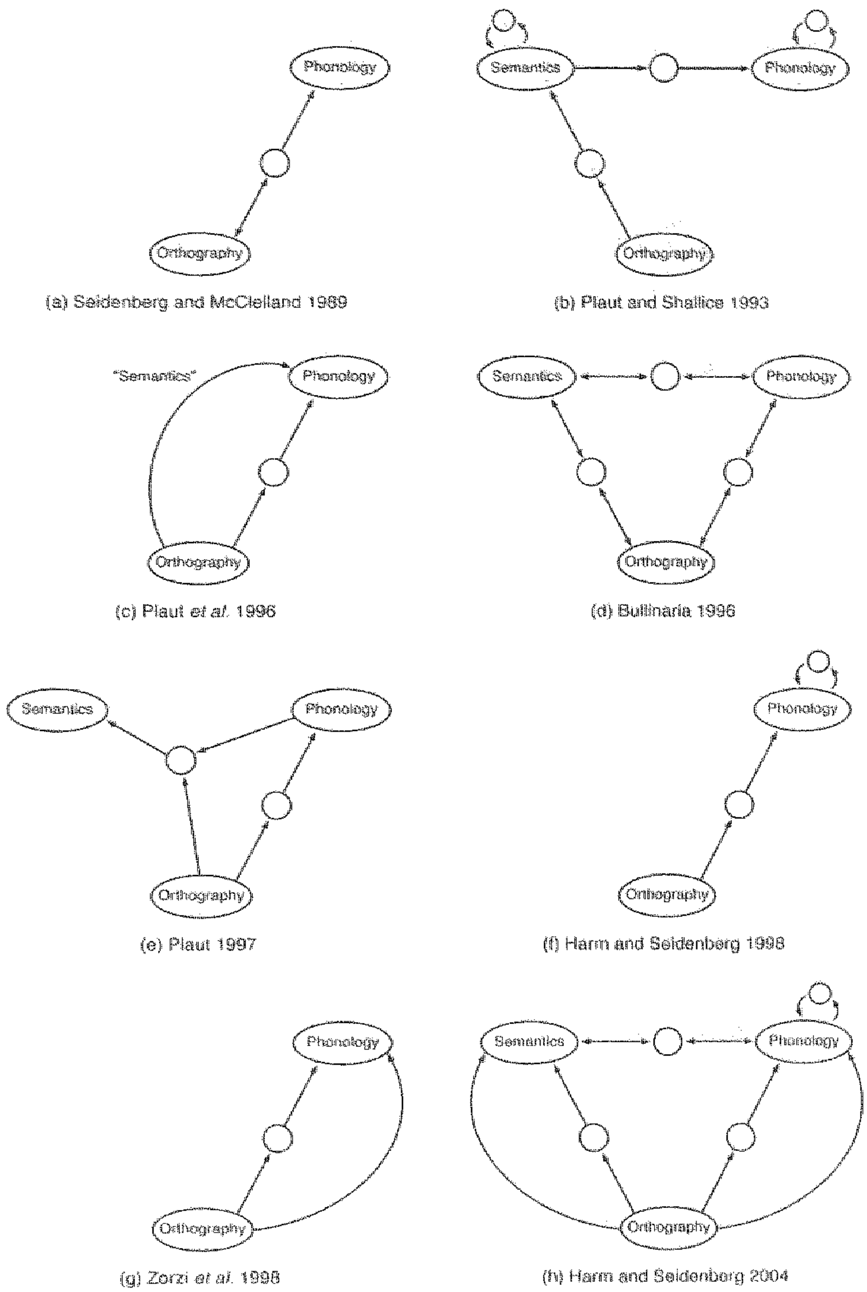

Figure 9.2. PDP models of word reading.

the idea of a multilayer network employing distributed representations with modifiable weights on connections) were employed in all of them. All the models have been concerned with processes involved in reading words, but they have emphasized different aspects of this skill: acquisition; skilled performance; developmental impairments (dyslexia); impairments following brain injury (acquired dyslexia); computation of pronunciations; computation of meanings; naming, lexical decision, and other tasks; bases of individual differences; and others. The models also differ with respect to many properties of the implementations, ranging from network architecture (e.g., number of layers, types of connectivity between layers) to training procedure (e.g., composition of the training corpus) to dependent measures (e.g., summed square error, settling time). Thus there has been a series of models which varied due to differences in focus with respect to behavioral phenomena (e.g., computation of pronunciation versus computation of meaning), and with respect to computational principles (e.g., interactive versus feedforward networks). The models also differ because of advances in understanding both the nature of the phenomena (e.g., the role of the orthography→semantics→phonology component in naming; factors that influence the division of labor between components of a model) and computational issues (e.g., how properties of network representations affect performance).

2.1 *Basic concepts*

The rationale behind this approach to understanding reading and other cognitive phenomena has been discussed extensively elsewhere (e.g., in the 1986 PDP volumes; in sources such as Plaut, 2005; Seidenberg, 1989, 1993; Seidenberg, 2005), and so only an overview is provided here. The goal of our research has been to develop a theoretical understanding of reading and how it is realized in the brain, with computational models providing the interface between the two. The idea is that it is not merely a nice thing to understand the brain bases of reading; rather, characteristics of reading (and other behaviors) arise from the ways in which reading is accomplished by the brain, and thus cannot be fully understood without reference to it.

This view is not universally accepted. The functional architecture style of theorizing assumed that neurobiology contributed very little to understanding the essential nature of the phenomena. The models of that era were not constrained by neurobiological facts (i.e., about how brains acquire, represent, and process information). Even though studies of brain-injured patients played a large role in model development, researchers focused on patterns of behavioral impairment rather than their neurobiological bases. (In fact I have a vivid memory of neurologist Norman Geschwind scolding researchers at an early 1980s conference for failing to report basic information about the nature of patient-subjects' brain injuries.) In part this functional stance was pragmatic, given the amount that was known about the neurobiology of cognition in the 1970s when the dual-route model and other functional style models were developed. However, this stance did not merely reflect lack of knowledge; in some circles there was also a philosophical commitment to theories of cognitive functions that abstracted away from neurobiology (e.g., Block and Fodor, 1972). Antireductionists such as Fodor (1999) have famously (if recklessly) declared their disinterest in exactly where above the neck cognitive functions are realized. This philosophy is largely retained in the DRC versions of the dual-route model, which are committed to computational explicitness but remain oriented to behavioral generalizations minimally constrained by biology.

In contrast, we assume that cognitive capacities such as reading are shaped by properties of the underlying neural substrate. Thus the methodology employed in our research involves formulating and evaluating computational principles that represent hypotheses about how neural activity gives rise to cognition. The question, then, is not merely what computations are involved or where they are realized in neural circuitry, but rather how the brain enables cognitive functions and why the brain arrived at particular solutions to computational problems. These principles are coupled with domain-specific considerations (e.g., facts and conditions specific to a task such as reading) in developing a theory of the phenomena. Hypotheses are tested by determining whether computational models that embody the proposed principles are, in fact, consistent with relevant behavioral and neurobiological data. Our understanding of these basic principles is partial and thus they are subject to revision as knowledge advances. In fact, there is feedback among all levels of analysis: behavioral, computational, and neurobiological. Discoveries at one level serve to constrain hypotheses at other levels, which in turn feedback on generating and testing hypotheses at other levels. Converging on the correct theory is literally a constraint satisfaction process (Seidenberg and MacDonald, 1999): The correct theory is that which satisfies constraints arising from the biological, computational, and behavioral levels.

The main principles that constitute the PDP framework are well-known: Behavior arises from the cooperative and competitive interactions among large networks of simple, neuron-like processing units; different types of information are represented by distributed patterns of activity over different groups of units, with similarity indexed by pattern overlap; knowledge is encoded as weights on connections between units; learning is the gradual adjustment of these weights based on the statistical structure among inputs and outputs; such networks may include internal, "hidden" representations that allow complex mappings to be learned; processing is constraint satisfaction, that is, the computed output best fits the constraints represented by the weights.

Several other important properties identified in connection with the reading models should also be noted.

1. The PDP principles on which they are based are general rather than domain-specific. Thus the same principles are thought to apply across perceptual, cognitive, and motor domains. That there are such general principles contrasts with the view that the brain has evolved many domain-specific subsystems (Pinker, 1997). One consequence of the domain-general view is that the mechanisms available for modeling are not introduced solely in response to the data from a particular domain, but are constrained to be consistent with applications in other domains. This means that the mechanisms in the reading models have some independent motivation; specifically, they are thought to reflect more general facts about human cognitive functions. This is particularly relevant to reading, a technology invented relatively recently in human history, making use of existing cognitive and perceptual capacities.
2. In the PDP approach, the models are a tool rather than the goal of the enterprise. Modeling is a means of exploring the validity and implications of a set of hypotheses about how cognitive processes are implemented in the brain. The goal is not the development of an individual model that can be taken as the account of some set of phenomena, based on a comfortable degree of fit to empirical data. Rather, the models facilitate converging on the correct set of explanatory principles, which are more general than any individual model, and on theories in particular substantive areas that employ these principles.

This orientation gives rise to two further characteristics of the research. One is that the failures and successes of the models are both sources of insight (McClelland et al., 1995). A classic illustration of this point is the Seidenberg and McClelland model's poor performance on nonwords, which definitely deviated from people's. The nonword generalization problem was soon traced to the imprecise way that phonological information was represented in the model (Plaut et al., 1996). This imprecision had little impact on the pronunciation of words, but affected nonword pronunciation because it is a harder task which requires recombining known elements in novel ways. Models with improved phonological representations yielded much better nonword performance (Harm and Seidenberg, 1999; Plaut et al., 1996). Thus the nonword problem "falsified" our original model, but not the theory it approximated. Moreover, this "failure" led to insights about how representations determine network behavior, to improved models, and to advances in understanding developmental dyslexia, which is associated with phonological impairments (Harm and Seidenberg, 1999). This pattern, in which the limitations of one model lead to deeper insights and improved next-generation models, is a positive aspect of the modeling methodology. However, it complicates the metric by which models are evaluated. A model could be a failure insofar as it did not capture some aspect of human behavior, but a success insofar as this limitation yielded greater insight about it.

A second consequence of the models as tools orientation is that the product of this research is a series of models that address different aspects of reading, which taken together advance the understanding of the phenomena. This approach to modeling is frustrating to some because there is no single simulation that constitutes the model of the domain. The models seem like a moving target: Seidenberg and McClelland's model was interesting but ultimately limited by its phonological representation; Plaut et al. (1996) largely fixed the phonological problem but introduced the idea that the orth→sem→phon pathway also contributes to pronunciation, something Seidenberg and McClelland had not considered. Harm and Seidenberg (1999) used yet another phonological representation and focused on developmental phenomena; Harm and Seidenberg (2004) implemented both orth→sem and orth→phon→sem parts of the triangle but focused on data concerning activation of meaning rather than pronunciation and so forth. Each model shares something with all of the others, namely the computational principles discussed above, but each model differs as well. Where, then, is the integrative model that puts the pieces all together?

The answer, of course, is that there is none. Achieving a complete, integrative model is an ill-conceived goal given the nature of the modeling methodology (particularly the need to limit the scope of a model in order to gain insights from it in finite time) and the goal of the enterprise, which, as in the rest of science, is the development of a general theory that *abstracts away* from details of the phenomena to reveal fundamental principles (Putnam, 1972). The models, as tools for exploring computational and empirical phenomena, change more rapidly than the theory, which is an abstraction away from individual models.

3. The PDP models are intrinsically about learning: They are literally systems that learn. Learning is central to PDP models for many reasons, the least of which is because algorithmic procedures happen to exist for training multilayer networks that employ distributed representations.

(i) The models are intended to address important questions about how people acquire information and represent it in ways that support complex behaviors. In reading (as in other areas such as language acquisition), we want to know how children acquire a skill given the nature of the task, the capacities they bring to the task, and the nature of their experience. Part of the explanation for why the reading system has the character that it does rests with facts about how it is learned.

(ii) Because they are systems that learn, the models provide a unified account of acquisition and skilled performance. The same principles govern both; children and adults represent different points on the developmental continuum represented by states of the model over training time.

(iii) The models address how knowledge representations develop. One consequence is that we avoid having to stipulate in advance how words are represented (e.g., whether there are codes for words, syllables, morphemes, etc.). Whether such representations exist is one of the basic questions and it is not addressed by building them into a model. The representations that the models develop are contingent on the input and output representations and general properties of the architecture such as numbers of units and layers. In fact the input and output representations are themselves learned, something an ideal network would also address. For the moment we can illustrate the general idea as in Figure 9.3. The fact that the models learn also obviates the need to wire a network by hand, as is done in many localist models. The problems associated with hand-wiring are discussed below in connection with DRC. The main problem is that hand-wiring promotes overfitting the results of particular studies. Moreover, the behavior of the models ends up depending on parameter settings that have no independent motivation or theoretical interpretation.

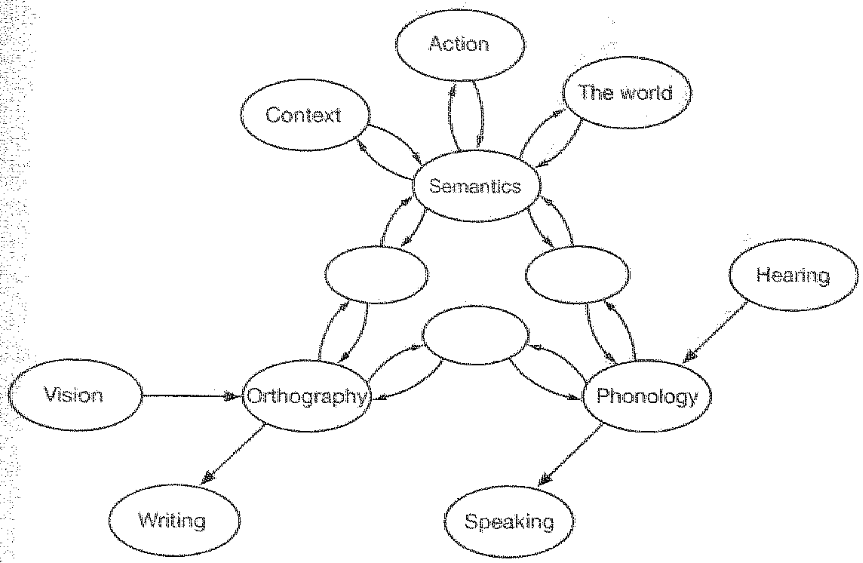

Figure 9.3. The orthographic, phonological, and semantic layers in the Seidenberg and McClelland model can be thought of as interlevel representations that result from sensation, perception, and action, in conjunction with constraints imposed by human biology.

> Yet these values are highly important insofar as model performance degrades when other values are used.

The emphasis on learning is one of the main characteristics that distinguishes the PDP and dual-route approaches. Although Coltheart et al. (1993) described an algorithm for deriving grapheme–phoneme correspondence rules, this procedure had little psychological plausibility and the rules that resulted were soon criticized for their oddness (Seidenberg et al., 1994). In later work, this algorithmic approach was abandoned and rules were added to DRC as needed to produce accurate output. The lack of a learning procedure for the lexical and nonlexical procedures in DRC is a serious limitation which extends beyond the need to hand-wire the rules themselves. The DRC models incorporate detailed procedures for applying rules to letter strings, which involve stipulations as to how the database of rules is searched, the order in which rules are applied, the size of the units over which rules apply, and other conditions. How the child would learn such scheduling procedures is also unknown. These gaps raise serious learnability questions: Could a child develop a dual-route system given children's experiences and the capacities and knowledge they bring to the task of learning to read? Coltheart and colleagues' view is apparently that because they model accounts for many aspects of skilled behavior, there must be some way for its components to be acquired. That inference is very weak, however, first because Coltheart et al. overstate the range of phenomena the model actually captures (see later in this chapter), and second because it is just as easy to imagine a model that does about as well as DRC2001 but could not be learned. The same issues do not arise in the PDP approach because learning is intrinsic to the enterprise.

My description of the PDP approach to reading presents an ambitious research agenda; what has actually been accomplished is of course limited in many ways. At this point in the study of reading, I think that having established the utility of the approach is more important than the details of any of the models, which are certain to change. Nonetheless, recognizing the limitations of the models is important and essential to further progress. This chapter is not long enough to tabulate all of these limitations and how they might be addressed in future work. Some examples, however: The goal is to understand reading and its brain bases, but the models have not as yet incorporated constraints arising from discoveries about the brain circuits that support reading. Such information (e.g., about the representation of semantic information in a distributed brain system closely tied to perception, action, and emotion) could be directly incorporated in future models via appropriate changes in the architecture. Many aspects of reading have simply not been addressed; for example, we have not seriously considered the visual front end – the recognition of letters and letter strings – or how context affects processing, particularly the ways in which the meanings of words shift as a function of the contexts in which they occur. The models are simplified at every turn; consider, for example, the differences between how the models learn and how people learn. The problem is not that the brain does not learn according to the backpropagation algorithm; "backpropagation of error" is not one of the foundational principles. More relevant is the fact that the models do not capture the range of experiences that occur in learning. A child is learning to read aloud. Sometimes she receives direct feedback about whether a word has been pronounced correctly or not; sometimes she can determine this herself (e.g., by using her computed pronunciation as the input to the comprehension system); sometimes she receives explicit feedback, other times just a reinforcement signal; sometimes there is no feedback except that provided by the child's own utterance. This variety of learning experiences could of course be incorporated in a model, and my own experience has been that model performance is not hurt by being more faithful to factors that govern human behavior. Whereas the models capture some important aspects of how children learn to read, they clearly do not address many others. The same could be said of almost every aspect of our models. It would be nice to have made further progress by now, but reading is a complex phenomenon involving most of human perception and cognition and most of the brain, and modeling is hard. Such is life.

3 The DRC models

The Dual-Route Cascade (DRC) model was described in a pair of papers by Coltheart and colleagues (1993; 2001), and its implications concerning reading acquisition, developmental and acquired forms of dyslexia, and other phenomena have been discussed in many additional publications (e.g., Coltheart, 2000; Coltheart, 2006; Jackson and Coltheart, 2001). The DRC framework includes mechanisms for computing both meanings and pronunciations; the implemented models exclusively focus on pronunciation.

Coltheart et al. (2001) offered extensive discussion of the origins of the DRC model and its fundamental assumptions. They linked their models to nineteenth-century diagram makers such as Lichtheim and emphasized continuity between the informal and computational versions of the dual-route model. The approach includes a commitment to a version of the modularity hypothesis (Coltheart, 1999; Fodor, 1983), to theorizing pitched at the level of the functional architecture, and to identifying the modules of the functional architecture primarily through studying brain-injured patients. However, the character of the DRC model is largely a consequence of two types of pretheoretical commitments. One concerns the goals of the modeling enterprise, the other the architecture of the model.

Coltheart et al. (2001) are explicit about their views of modeling. They emphasize the data-driven character of their modeling, endorsing Grainger and Jacobs' (1998, p. 24) view that "in developing algorithmic models of cognitive phenomena, the major source of constraint is currently provided by human behavioral data." They also view models as cumulative: Each model improves upon the previous one by adding to the phenomena covered by a previous version. Thus, the 2001 version of the DRC is said to account for the same facts as the 1999 version but also many others. The models are said to be nested with respect to data coverage. Again they quote Grainger and Jacobs (1998): "[I]n other sciences it is standard practice that a new model accounts for the crucial effects accounted for by the previous generation of the same or competing models."

This gives rise to an approach in which fidelity to behavioral data is the principal criterion for evaluating models. Models are designed to recreate data patterns, and a model is valid unless disconfirmed by making an incorrect prediction. So, for example, in justifying their use of an interactive activation model as a component of DRC, the authors note that the McClelland and Rumelhart IA model had not been refuted by any behavioral data; thus there was no empirical reason to abandon it. The strength of the 2001 version of DRC, then, is the fact that it addresses more than twenty different phenomena. The breadth of the data coverage led Coltheart and colleagues to conclude that "the DRC model is the most successful of the existing computational models of reading" (p. 204).

The second type of commitment is architectural. Coltheart et al. (1993; 2001) retain the central dogma of the dual-route approach, that two mechanisms are required for words and nonwords: a nonlexical route consisting of rules mapping graphemes onto phonemes, and a lexical route involving word-specific knowledge. On the nonlexical side, the implemented models spell out details about what the rules are, expanding the notion of rule to include multiletter graphemes, and adjudicating the order in which rules are applied. They also incorporate assumptions about the timing of rule application, which affects the extent to which output from the lexical and nonlexical routes is combined in generating a pronunciation. On the lexical side, they incorporate a version of the McClelland and Rumelhart (1981) interactive activation model, though with rather different parameter settings than in the original work.

The central dogma reflects a strongly held intuition about the nature of human knowledge and the behaviors it supports, namely that they are rule-governed. Language provides a quintessential example, but the same intuition extends to other areas such as reasoning and social behavior. People being complex organisms, the rules that govern their behavior often have exceptions. Thus the mind/brain has developed other mechanisms for dealing with such deviations (e.g., memorization). Perhaps these intuitions are strong because prior to the development of connectionist models there were few if any alternatives. Perhaps it persists because it meets the requirements of a good folk psychological theory: It is easy to understand and at a very general level seems to be correct. Of course it is trivially true that *any* body of knowledge or behavior can be described as rule-governed if there is a separate mechanism to handle cases that violate the rules. The number of phenomena to be explained (rule-governed cases and exceptions) exactly matches the number of explanatory mechanisms (a rule component, an exception component). Still, the intuition seems to be a strong one, and certainly more accessible than the idea of a multilayer neural network using distributed representations that encode the statistical, probabilistic mappings between spelling and sound. It could also be correct, which of course is the force of Coltheart et al.'s accounting of the many phenomena that DRC simulates.

Given these assumptions, what type of computational model results? Seidenberg and Plaut (2006) discuss this issue. Coltheart et al. (2001) term their approach "Old Cognitivism," identifying it within a longstanding tradition of fitting models to behavioral data. However, they did not discuss longstanding critiques of this approach. The basic problem is that the strategy of accounting for the most data possible is not itself sufficiently powerful to converge on a satisfactory theory. Such models are built out of an ad hoc collection of formalisms such as rules and buffers and decision makers. Elements of the model are configured in response to empirical data; that is the essence of the methodology. This approach does not yield robust theoretical generalizations that explain target phenomena.

As Seidenberg and Plaut argue, the evidence for this conclusion is the fact that the performance of the 2001 version of DRC is closely tied to the specific studies that were simulated. In most cases the authors chose a single study to demonstrate DRC's coverage of a given phenomenon. However, the model fails when tested against other studies of the same type. For example, many studies have reported that frequency and regularity interact in reading aloud (see Seidenberg, 1995 for review). Coltheart et al. simulated the results of one of these studies, Paap and Noel (1989). There are two problems. First, the simulation produces a frequency by regularity interaction, but it is not of the same form as that observed in the behavioral study. Second, the model fails to reproduce the interaction when tested on the stimuli used in other classic studies (e.g., Seidenberg et al., 1984; Taraban and McClelland, 1987). This pattern is repeated for most of the major phenomena that DRC2001 simulates. The best that can be said is that the model's coverage of the data is broad but shallow. The worst is that failures to correctly simulate effects of many studies disconfirm it. Seen in this light, Coltheart et al.'s claim to have tied off a long list of phenomena seems an overstatement of considerable proportion.

Fitting models to a broad range of data is very difficult. DRC does not have a learning procedure and so all of the parameters of the model (Table 1, p. 218, in the 2001 article lists thirty-one, which does not include seven additional parameters for the lexical decision task) must be set by painstakingly searching the parameter space for the values that will best fit the broadest range of studies (i.e., hand-wiring). Coltheart et al. were able to find a set of parameters that correctly simulated individual studies; however these values yield incorrect results when other studies of the same effects are considered. DRC's problem is *overfitting*: The model is too closely tied to the specific studies that were simulated and thus it fails to generalize.

How could this overfitting problem be avoided? One thought is that a model as complex as DRC, with many parameters that can be independently manipulated, should fit any pattern of data. However, that does not turn out to be true. To the contrary, DRC's performance is closely tied to the particular parameter values that were chosen. Changes to the parameters tend to produce worse behavior, not better.[2] Another possibility is that the model fails because it does not operate according to the principles that underlie the corresponding human behavior. This seems likely but ironic given Coltheart et al.'s explicit disavowal of principle-based approaches to modeling.

The validity of these conclusions is also supported by developments since the publication of the 2001 version of DRC. Coltheart and colleagues have presented additional data and arguments favoring the DRC model, primarily focused on atypical cases of developmental or acquired dyslexia with patterns of impaired reading that are said to contradict one or another aspect of the PDP theory (e.g., Blazely, Coltheart, and Casey, 2005; Coltheart, 2006). The interpretation of such "outlier" cases is highly controversial (see Plaut, 1995 and Woollams et al., 2006 for discussion). The more important point is that Coltheart and colleagues' analyses of these cases are not tied to their implemented model. They wish to argue that certain patterns of behavioral impairment are consistent with their model (and not the PDP approach); however, these arguments are not coupled to simulations of the cases in question.

I don't think this is surprising. According to the analysis presented earlier in this chapter, the successes of the 2001 version of DRC are so tightly bound to particular parameter values and other implementation-specific features that it *should* be difficult to extend it to other phenomena, including these case studies. As a result, the most recent arguments on behalf of the dual-route approach

2 Exploration of the model suggests that changes to the parameters tend to degrade performance. However, I have no basis for excluding the possibility that slightly better fits might be achieved by searching the parameter space further. Rather, I am suggesting that the problem is deeper than mere parameter search.

have the same form as in the precomputational modeling era discussed in the first section of this chapter. It is almost as though the models (and the insights about the need for such models in order to understand normal and disordered behavior) had never occurred.

4 Hybrid models

Two prominent hybrid models purport to incorporate the best aspects of the dual-route and PDP approaches. The first, by Zorzi, Houghton, and Butterworth (1998), was an SM89-style model of the orthography–phonology mapping that incorporated an additional set of connections between the orthographic and phonological units. Zorzi et al. claimed that the result was a "dual-route connectionist model." The connections between the orthographic and phonological (input and output) units were said to correspond to the sublexical route, whereas the orthography-hidden-phonological pathway corresponded to the lexical route. Thus, the model adhered to the central dogma, but used connectionist machinery. As Harm and Seidenberg (2004) determined, however, the Zorzi et al. model did not create a division of labor such that regulars were handled by one pathway and exceptions by the other. The direct O–P pathway, acting on its own (i.e., with the rest of the model turned off or "lesioned"), read regular words accurately as well as a small percentage of exception words (the latter indicating that it was not strictly nonlexical). With the direct O–P pathway disabled, however, the standard orthography-hidden-phonology pathway – the putative lexical route – read exception words highly inaccurately. In their model, reading exception words requires input from both pathways. This is similar to the account of exception word reading in the triangle framework and quite different from the dual-route approach.

Perry, Ziegler, and Zorzi (2007) created a different type of hybrid. They remain committed to the central dogma that distinct lexical and sublexical procedures are necessary; however they abandoned DRC's commitment to the notion that the sublexical route consists of grapheme-phoneme mapping rules. Instead, they employed a PDP-style model for the nonlexical route. As in DRC, the lexical route is an interactive activation model. This route is thought necessary in order to account for data about lexical decisions.

The CDP+ model has plusses and minuses. On the positive side, Perry et al. incorporated an SM89-style O–P route in recognition of the fact that models like DRC cannot generate consistency effects. They report positive results for a close simulation of a study by Jared (2002). This move further validates the approach we initiated in 1989. But in other ways the Perry et al. model fares less well. The authors emphasize the importance of capturing "benchmark" effects that have the potential to differentiate between alternative theories. Consistency effects are one such benchmark, regularity effects are another, and so on. Targeting theoretically important phenomena is a fine idea and one that the triangle approach instantiated; the question is whether a given model captures them in a principled way or not. Here the Perry et al. model falls into the same trap as DRC2001: It simulates individual studies (such as one experiment out of Jared's 2002 article) but fails when tested on other studies of the same phenomena. It too overfits the data, but this matters less for CDP+, the authors believe, because they have been selective in choosing which studies to simulate. For example, the Jared experiment is taken as the gold standard for consistency effects and so that is the study they simulate. I share the view that models should aim to account for prominent benchmark phenomena but disagree with their reliance on gold standard experiments. Jared is a fine study, but other experiments that yielded the same effects cannot be dismissed out of hand. The methods in this area of research are weak enough that people have tended to rely on replication with different materials and subjects, across different labs, with the hope of gaining strong converging evidence for a phenomenon. It seems a mistake to use a different method in assessing models of these phenomena. The Perry et al. model also misses many other effects that many people consider benchmarks, including the frequency x regularity interaction in studies by Taraban and McClelland (1987), myself, and others, and the generation of correct pronunciations for complex nonwords.[3]

Perry et al.'s other major theoretical claim is that a localist-style interactive activation model is necessary in order to account for lexical decision phenomena. They therefore incorporated the mechanism from DRC2001 unchanged. This account of lexical decision is highly engineered and performance is brittle, degrading with changes to the seven task-specific parameters. The model captures a narrow range of LD data and again it is easy to find studies on which it fails (e.g., Ziegler and Perry's 1998 study of orthographic versus body neighborhood effects). A flawed account of the lexical decision task does not provide strong motivation for retaining DRC's version of the lexical route. Having incorporated an SM89-style orth–phon system, my suggestion would be to take the next step, which would be to incorporate the orth-sem-phon pathway (as in Harm and Seidenberg, 2004). Lexical decisions might then be based on different patterns of activation produced by words and nonwords over the different types of units, an idea (Seidenberg and McClelland, 1989) that has not been fully explored.

3 Perry et al.'s model errs in generating pronunciations for difficult nonwords. They suggest that such data are outside the scope of their model because subjects use metalinguistic problem solving strategies to pronounce them. This is an odd turn of events, since the principal complaint about the SM89 model was its poor performance on such nonwords. By Besner et al.'s (1990) criteria, Perry et al.'s model is in the same boat we were in 1989.

This was the approach to lexical decision we proposed in SM89. In simulating LD, the model made more errors than people did, but, importantly, it contained no semantics. People do occasionally consult meaning in order to decide if a letter string is a word or not. The idea that lexical decisions are made by comparing the patterns of activation over orthographic, phonological, and semantic representations produced by words versus nonwords remains a viable though underexplored approach to understanding the task.

5 Conclusions

The competition between the dual-route model and the alternative PDP approach has been highly productive in many ways. It stimulated a large amount of excellent empirical research on many aspects of word and nonword reading. It led to more explicit, mechanistic accounts of behavior. The models have grown in sophistication and scope. However, as time has progressed the differences between the approaches have become clearer. The PDP models address a narrower range of phenomena, but aspire to considerable depth insofar as the explanations for the phenomena (e.g., consistency effects, the relationship between frequency and consistency, nonword pronunciation, division of labor between components of a model) are meant to derive from the principles that govern the approach rather than details of how particular models are implemented. In fact the models that we have published could have achieved better fits to particular data sets by manipulating factors such as the composition of the training set or parameters of the learning algorithm. One of the main points of this chapter is that such manipulations are self-defeating unless they are motivated by independently known facts (e.g., about people's experience or how they learn). One can win the best fits battle and lose the correct theory war.

The DRC models' success and failures arise from their creators' data fitting orientation. In advocating this approach, I think that Coltheart and colleagues have misread the history of cognitive science and cognitive modeling. Research has not progressed in the cumulative, nested manner described by Grainger and Jacobs (1998) and endorsed by Coltheart et al. (2001). Researchers who have utilized computational models have tended to develop a series of models that share common principles but differ in detail and address overlapping but nonidentical phenomena (e.g., Roger Ratcliff's diffusion models; John Anderson's ACT models). The same could be said of DRC. The 1993 version of DRC did not start by replicating the results reported by McClelland and

Rumelhart and then extending the model to other phenomena. The two implemented versions of DRC share many properties but also differ in detail; the phenomena they address only partially overlap. In this respect the situation is not very different from the PDP models. I doubt if any science proceeds in a strictly cumulative manner, but that is a matter for philosophers and historians of science to decide.

What about the PDP models? I think we have made considerable progress with respect to those reading phenomena that have been addressed, which include ones that are highly revealing about the basic character of the system (e.g., consistency effects). As noted, however, many aspects of reading remain to be addressed within this framework. The approach has limitations, but of a different character than those of DRC. The main problem seems to be that many researchers think that, despite the talk about principles, the models do little to clarify the bases of complex behavior. These concerns were voiced by McCloskey (1991), who noted that in connectionist models, the explanations for behavior tended to be buried in a mass of units and connections. I think McCloskey was wrong (Seidenberg, 1993), but perhaps this is a minority opinion. Clearly the methodology that underlies DRC can also be applied in developing PDP models, with similar results: creating models to fit the results of particular studies, resulting in brittle models whose performance is highly dependent on unmotivated implementational details. Nor have PDP modelers been exempt from premature declarations of victory. PDP modeling is an approach that can be used in different ways by different people; some applications will be more helpful than others. I think the more important point, however, is that the PDP modelers have apparently not done an adequate job of explaining the philosophy and the approach, the goals of the enterprise, the general principles that underlie the models, and the reasons why models behave the way they do. My own feeling is that these issues have already been addressed in the literature but perhaps not successfully.

The history of computational models of reading suggest several conclusions that could profitably inform future research: that models are a tool but theories are the goal; that a principled approach is more likely to yield progress than data fitting; that progress has to be assessed in terms of a longer term trajectory toward a goal, rather than comparisons of specific models that only represent points on that trajectory. I hope that researchers continue to recognize the enormous opportunities for making further progress in understanding reading, computation, and the brain.

References

Adams, M. J. (1979). Models of word recognition. *Cognitive Psychology*, 11, 133–76.

Baron, J. & Strawson, C. (1976). Use of orthographic and word-specific knowledge in reading words aloud. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 207–14.

Besner, D., Twilley, L., McCann, R., & Seergobin, K. (1990). On the connection between connectionism and data: Are a few words necessary? *Psychological Review*, 97, 432–46.

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

Block, N. & Fodor, J. (1972). What psychological states are not. *Philosophical Review*, 81–92.

Caplan, D. (1988). On the role of group studies in neuropsychological and pathopsychological research. *Cognitive Neuropsychology*, 5, 535–47.

Caramazza, A. (1986). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain & Cognition*, 5, 41–66.

Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, 7, 161–89.

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

(1999). Modularity and cognition. *Trends in Cognitive Sciences*, 3, 115–20.

(2000). Dual routes from print to speech and dual routes from print to meaning: Some theoretical issues. In Kennedy, A., Radach, R., Pynte, J., & Heller, D. (Eds.) *Reading as a Perceptual Process*. Oxford: Elsevier.

(2006). Acquired dyslexias and the computational modeling of reading. *Cognitive Neuropsychology*.

Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual-route and parallel distributed processing approaches. *Psychological Review*, 100, 589–608.

Coltheart, M., Patterson, K. E., & Marshall, J. C. (Eds.) (1980). *Deep dyslexia*. London: Routledge & Kegan Paul.

Coltheart, M., Rastle, K., Perry, C., & Langdon, R. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–56.

Dunn, J. C. & Kirsner, K. (Eds.) (2003). Forum on "What Can We Infer From Double Dissociations," *Cortex*, 39, 129–202.

Farah, M. J. (1994). Neuropsychological inference with an interactive brain: A critique of the "locality" assumption. *Behavioral and Brain Sciences*, 17, 43–104.

Fodor, J. A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.

(1999). Let your brain alone. *London Review of Books*, September 30.

Glushko, R. J. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, 5, 674–91.

Grainger, J., & Jacobs, A. M. (1994). Orthographic processing in visual word recognition. *Psychological Review*, 103, 518–65.

(1998). On localist connectionism and psychological science. In Grainger, J. & Jacobs, A. M. (Eds.) *Localist connectionist approaches to human cognition* (pp. 1–38). Mahwah, NJ: Erlbaum.

Harm, M. & Seidenberg, M. S. (1999). Reading acquisition, phonology, and dyslexia: Insights from a connectionist model. *Psychological Review*, 106, 491–528.

(2004). Computing the meanings of words in reading: Division of labor between visual and phonological processes. *Psychological Review*, 111, 662–720.

Jackson, N. E. & Coltheart, M. (2001). *Routes to Reading Success and Failure*. New York: Psychology Press.

Jared, D. (2002). Spelling sound consistency and regularity effects in word naming. *Journal of Memory and Language*, 46, 723–50.

Juola, P. & Plunkett, K. (2000). Why double dissociations don't mean much. In Cohen, G., Johnston, R., & Plunkett, K. (Eds.) *Exploring cognition: Damaged brains and neural networks* (pp. 111–72). New York: Psychology Press.

Kay, J. & Marcel, A. (1981). One process, not two, in reading aloud: Lexical analogies do the work of non-lexical rules. *Quarterly Journal of Experimental Psychology*, 33A, 397–413.

Luce, R. D. (1959). *Individual Choice Behavior*. New York: Wiley.

Marshall, J. C. & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research*, 2, 175–9.

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375–407.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419–57.

McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. *Psychological Science*, 2, 387–95.

Miozzo, M. (2003). On the processing of regular and irregular forms of verbs and nouns: evidence from neuropsychology. *Cognition*, 87, 101–27.

Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165–78.

Newell, A. & Simon, H. (1963). GPS, a program that simulates human thought. In Feigenbaum, E. A. & Feldman, J. (Eds.) *Computers and Thought*. New York: McGraw Hill.

Paap, K. & Noel, R. W. (1989). Dual-route models of print to sound: Still a good horse race. *Psychological Research*, 53, 13–24.

Patterson, K. E., Marshall, J. C., & Coltheart, M. (Eds.) (1985). *Surface dyslexia*. London: Erlbaum.

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

Pinker, S. (1997). *How the Mind Works*. New York: Norton.

Plaut, D. C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *J. Clinical and Exp. Neuropsychology*, 17, 291–321.

(2005). Connectionist approaches to reading. In Snowling, M. J. & Hulme, C. (Eds.) *The Science of Reading: A Handbook* (pp. 24–38). Blackwell: Oxford.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasiregular domains. *Psychological Review*, 103, 56–115.

Putnam, H. (1972). Psychology and reduction. *Cognition*, 2, 131–46.

Rapp, B. & Caramazza, A. (2002). Selective difficulties with spoken nouns and written verbs: A single case study. *Journal of Neurolinguistics*, 15, 373–402.

(1993). On the distinction between deficits of access and deficits of storage: A question of theory. *Cognitive Neuropsychology*, 10, 113–41.

Reicher, G. M. (1969). Perceptual recognition as a function of meaningfulness of the stimulus material. *Journal of Experimental Psychology*, 81, 274–80.

Riddoch, M. J., Humphreys, G. W., Coltheart, M., & Funnell, E. (1988). Semantic systems or system? Neuropsychological evidence re-examined. *Cognitive Neuropsychology*, 5, 3–25.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In Rumelhart, D. E. & McClelland J. L. (Eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition (vol. 1)*. Cambridge, MA: MIT Press.

Seidenberg, M. S. (1989). Word recognition and naming: A computational model and its implications. In Marslen-Wilson, W. D. (Ed.) *Lexical Representation and Process*. MIT Press.

(1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology*, 5, 403–26.

(1993). A connectionist modeling approach to word recognition and dyslexia. *Psychological Science*, 4, 299–304.

(1995). Visual word recognition. In Miller, J. L. & Eimas, P. D. (Eds.) *Handbook of Perception & Cognition, Volume 11, Speech, Language & Communication*. San Diego: Academic Press.

(2005). Connectionist models of reading. *Current Directions in Psychological Science*, 14, 238–42.

Seidenberg, M. S. & MacDonald, M. C. (1999). A probabilistic constraints approach to language acquisition and processing. *Cognitive Science*, 23, 569–88.

Seidenberg, M. S. & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523–68.

Seidenberg, M. S., & Plaut, D. C. (2006). Progress in understanding word reading: Data fitting vs. theory building. In Andrews, S. (Ed.) *From Ink Marks To Ideas: Current Issues In Lexical Processing*. Hove, UK: Psychology Press.

Seidenberg, M. S., Plaut, D., Petersen, A., McClelland, J. L., & McRae, K. (1994). Nonword naming and models of word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1177–96.

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 Behavior*, 23, 383–404.

Shallice, T. (1993). Multiple semantics: Whose confusions? *Cognitive Neuropsychology*, 10, 251–61.

Taraban, R. & McClelland, J. L. (1987). Conspiracy effects in word recognition. *Journal of Memory and Language*, 26, 608–31.

Van Orden, G., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove? *Cognitive Science*, 25, 111–72.

Wagenmakers, E. J. M., Steyvers, M., Raaijmakers, J. G. W., Shiffrin, R. M., van Rijn, H., & Zeelenberg, R. (2004). A model for evidence accumulation in the lexical decision task. *Cognitive Psychology*, 48, 332–67.

Woollams, A. M., Lambon Ralph, M. A., Plaut, D. C., & Patterson, K. E. (2006). SD-Squared: On the association between semantic dementia and surface dyslexia. *Psychological Review*.

Ziegler, J. C. & Perry, C. (1998). No more problems in Coltheart's neighborhood: resolving neighborhood conflicts in the lexical decision task. *Cognition, 68(2)*, 53–62.

Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception & Performance*, 24, 1131–61.