








PART I

Word Recognition Processes in Reading





Editorial Part I

Word recognition is the foundation of reading; all other processes are dependent on it. If word recognition processes do not operate fluently and efficiently, reading will be at best highly inefficient. The study of word recognition processes is one of the oldest areas of research in the whole of experimental psychology (Cattell, 1886). The chapters in this section of the Handbook present an overview of current theories, methods, and findings in the study of word recognition processes in reading.

What do we mean by recognition here? Recognition involves accessing information stored in memory. In the case of visual word recognition this typically involves retrieving information about a word's spoken form and meaning from its printed form. The first two chapters, by Coltheart and Plaut, outline the two most influential theoretical frameworks for studies of visual word recognition.

Coltheart outlines the history and evolution of dual-route models of reading *aloud* (i.e., how the pronunciation of a printed word is generated). These dual-route models posit that there are two routes from print to speech: a lexical and nonlexical route. Broadly the lexical route involves looking up the pronunciation of a word stored in a lexicon or mental dictionary. In contrast, the nonlexical route involves translating the graphemes (letters or letter groups) into phonemes and assembling the pronunciation of a word from this sequence of phonemes. Such a process should work just as well for nonwords as for words, just so long as the word follows the spelling pattern of the language (a nonlexical reading of YACHT, will not yield the pronunciation for a kind of boat with a sail on it). This idea is embodied in an explicit computational model (the DRC model) that Coltheart describes in detail. It may be worth emphasizing that this highly influential model is a model of how adults read aloud; it is not concerned with how the knowledge allowing this to happen is acquired. A major focus of the model is how different disorders of reading aloud, which arise after brain damage in adults, can be accounted for.

Plaut gives an overview of a different class of models of reading aloud that employ connectionist architectures (models that learn to pronounce words by training associa-

4 *Editorial Part I*

tions between distributed representations of orthography and phonology). One particularly influential model of this type is the so-called triangle model (Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989). This model abandons the distinction between a lexical and nonlexical procedure for translating visual words into pronunciations; instead the same mechanism is used to convert words and nonwords into pronunciations, based on patterns of connections between orthographic inputs and phonological outputs. One other critical difference between the triangle model and the DRC model is that the triangle model explicitly embodies a learning procedure and thus can be considered a model of both adult reading and reading development. It is clear that these are very different conceptions of how the mind reads single words. Both approaches deal with a wide range of evidence. Arguably, the DRC model is more successful in dealing with the detailed form of reading impairments observed after brain damage in adults, while the ability to think about development and adult performance together in the triangle model is a considerable attraction. There is no doubt that differences between these models will be a source of intense interest in the coming years.

Lupker's chapter moves on to review a huge body of experimental evidence concerned with how adults recognize printed words. Many of these experiments investigate what is a remarkably rapid and accurate process in most adults, by measuring reaction time, or by impairing performance by using masking (preventing participants from seeing a word clearly by superimposing another stimulus immediately after the word has been presented). Any complete model of word recognition ultimately will have many phenomena from such experiments to explain. These include the fact that people perceive letters more efficiently when they are embedded in words, that high-frequency (i.e., more familiar) words are recognized easier than less familiar words, and that recognition of words is influenced by previously presented words (seeing a prior word that is related in form or meaning helps us to recognize a word that follows it). One conclusion that emerges powerfully from Lupker's review is the need for interactive models in which activation of orthographic and phonological information reciprocally influence each other. This is an issue that Van Orden and Kloos take up in detail, presenting a wealth of evidence that converges on the idea that there is intimate and perpetual interaction between representations of orthography and phonology (spelling and sound) during the process of recognizing a printed word.

Moving on from the recognition of isolated words, Rayner, Juhasz, and Pollatsek discuss eye movements in reading. Eye movements provide a fascinating window on how word recognition processes operate in the more natural context of reading continuous text. It appears that the pattern of eye movements in reading is heavily influenced by the cognitive processes subserving both word recognition and text comprehension. The majority of words in text are directly fixated (usually somewhere in the first half of the word). For readers of English the area of text processed during a fixation (the perceptual span) is about 3 or 4 letters to the left of fixation and some 14 or 15 letters to the right of fixation. This limit seems to be a basic one determined by acuity limitations, and useful information about letter identity is extracted only from a smaller area, perhaps 7 or 8 letters to the right of the fixation point. It appears that only short, frequent, or highly predictable words are identified prior to being fixated (so that they can be skipped). However, partial information (about a word's orthography and phonology but typically

not its meaning) about the word following the fixation point often is extracted and combined with information subsequently extracted when the word is directly fixated. These studies are consistent with the view that the speed and efficiency of word recognition processes (as well as higher-level text-based processes) place fundamental constraints on how quickly even skilled readers read text.

Arguably the central question in the study of word recognition in reading is the role of phonology. All of the chapters in Part I address this issue explicitly. It appears that a consensus has been reached: phonological coding is central to word recognition, though opinions are divided on many details of how phonology is accessed and its possible importance in providing access to semantic information.

1

Modeling Reading: The Dual-Route Approach

Max Coltheart

Reading is information-processing: transforming print to speech, or print to meaning. Anyone who has successfully learned to read has acquired a mental information-processing system that can accomplish such transformations. If we are to understand reading, we will have to understand the nature of that system. What are its individual information-processing components? What are the pathways of communication between these components?

Most research on reading since 1970 has investigated reading aloud and so sought to learn about the parts of the reading system that are particularly involved in transforming print to speech. A broad theoretical consensus has been reached: whether theories are connectionist (e.g., Seidenberg & McClelland, 1989; Plaut, this volume) or nonconnectionist (e.g., Coltheart, Curtis, Atkins & Haller, 1993), it is agreed that within the reading system there are two different procedures accomplishing this transformation – there are dual routes from print to speech. (The distinction between connectionist and nonconnectionist theories of cognition is discussed later in this chapter.)

In the Beginning . . .

The dual-route conception of reading seems first to have been enunciated by de Saussure (1922; translated 1983, p. 34):

there is also the question of reading. We read in two ways; the new or unknown word is scanned letter after letter, but a common or familiar word is taken in at a glance, without bothering about the individual letters: its visual shape functions like an ideogram.

However, it was not until the 1970s that this conception achieved wide currency. A clear and explicit expression of the dual-route idea was offered by Forster and Chambers (1973):

The pronunciation of a visually presented word involves assigning to a sequence of letters some kind of acoustic or articulatory coding. There are presumably two alternative ways in which this coding can be assigned. First, the pronunciation could be computed by application of a set of grapheme–phoneme rules, or letter–sound correspondence rules. This coding can be carried out independently of any consideration of the meaning or familiarity of the letter sequence, as in the pronunciation of previously unencountered sequences, such as *fitch*, *mantiness* and *streep*. Alternatively, the pronunciation may be determined by searching long-term memory for stored information about how to pronounce familiar letter sequences, obtaining the necessary information by a direct dictionary look-up, instead of rule application. Obviously, this procedure would work only for familiar words. (Forster & Chambers, 1973, p. 627)

Subjects always begin computing pronunciations from scratch at the same time as they begin lexical search. Whichever process is completed first controls the output generated. (Forster & Chambers, 1973, p. 632)

In the same year, Marshall and Newcombe (1973) advanced a similar idea within a box-and arrow diagram. The text of their paper indicates that one of the routes in that model consists of reading “via putative grapheme–phoneme correspondence rules” (Marshall & Newcombe, 1973, p. 191). Since the other route in the model they proposed involves reading via semantics, and is thus available only for familiar words, their conception would seem to have been exactly the same as that of Forster and Chambers (1973).

This idea spread rapidly:

We can . . . distinguish between an orthographic mechanism, which makes use of such general and productive relationships between letter patterns and sounds as exist, and a lexical mechanism, which relies instead upon specific knowledge of pronunciations of particular words or morphemes, that is, a lexicon of pronunciations (if not meanings as well). (Baron & Strawson, 1976, p. 386)

It seems that both of the mechanisms we have suggested, the orthographic and lexical mechanisms, are used for pronouncing printed words. (Baron & Strawson, 1976, p. 391)

Naming can be accomplished either by orthographic-phonemic translation, or by reference to the internal lexicon. (Frederiksen & Kroll, 1976, p. 378)

In these first explications of the dual route idea, a contrast was typically drawn between words (which can be read by the lexical route) and nonwords (which cannot, and so require the nonlexical route). Baron and Strawson (1976) were the first to see that, within the context of dual-route models, this is not quite the right contrast to be making (at least for English):

The main idea behind Experiment 1 was to compare the times taken to read three different kinds of stimuli: (a) regular words, which follow the “rules” of English orthography, (b) exception words, which break these rules, and (c) nonsense words, which can only be pronounced by the rules, since they are not words. (Baron & Strawson, 1976, p. 387)

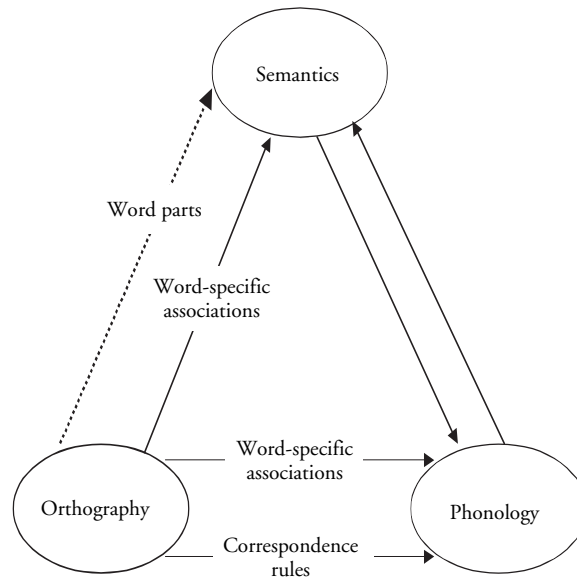


Figure 1.1 An architecture of the reading system (redrawn from Baron, 1977).

Baron (1977) was the first to express these ideas in a completely explicit box-and-arrow model of reading, which is shown in figure 1.1. This model has some remarkably modern features: for example, it has a lexical-nonsemantic route for reading aloud (a route that is available only for words yet does not proceed via the semantic system) and it envisages the possibility of a route from orthography to semantics that uses word parts (Baron had in mind prefixes and suffixes here) as well as one that uses whole words.

Even more importantly, the diagram in figure 1.1 involves two different uses of the dual-route conception. The work previously cited in this chapter all concerned a dual-route account of reading aloud; but Baron's model also offered a dual-route account of reading comprehension:

we may get from print to meaning either directly – as when we use pictures or maps, and possibly when we read a sentence like *I saw the son* – or indirectly, through sound, as when we first read a word we have only heard before. (Baron, 1977, p. 176)

Two different strategies are available to readers of English for identifying a printed word. The phonemic strategy involves first translating the word into a full phonemic (auditory and/or articulatory) representation, and then using this representation to retrieve the meaning of the word. This second step relies on the same knowledge used in identifying words in spoken language. This strategy must be used when we encounter for the first time a word we have heard but not seen. The visual strategy involves using the visual information itself (or possibly some derivative of it which is not formally equivalent to overt pronunciation) to retrieve the meaning. It must be used to distinguish homophones when the context is insufficient, for example, in the sentence, “Give me a pair (pear).” (Baron & McKillop, 1975, p. 91)

The dual-route theory of reading aloud and the dual-route theory of reading comprehension are logically independent: the correctness of one says nothing about the correctness of the other. Further discussion of these two dual-route theories may be found in Coltheart (2000). The present chapter considers just the dual-route approach to reading aloud.

A final point worth making re Baron's chapter has to do with the analogy he used to illustrate why two routes might be better than one (even when one is imperfect – the nonlexical route with irregular words, for example):

A third – and to me most satisfying – explanation of the use of the indirect path . . . is that it is used in parallel with the direct path. If this is the case, we can expect it to be useful even if it is usually slower than the direct path in providing information about meaning. If we imagine the two paths as hoses that can be used to fill up a bucket with information about meaning, we can see that addition of a second hose can speed up filling the bucket even if it provides less water than the first. (Baron, 1977, p. 203)

An analogy commonly used to describe the relationship between the two routes in dual-route models has been the horse race: the lexical and nonlexical routes race, and whichever finishes first is responsible for output. But this analogy is wrong. In the reading aloud of irregular words, on those occasions where the nonlexical route wins, according to the horse race analogy the response will be wrong: it will be a regularization error. But what is typically seen in experiments on the regularity effect in reading aloud is that responses to irregular words are correct but slow. The horse race analogy cannot capture that typical result, whereas Baron's hose-and-bucket analogy can. The latter analogy is equally apt in the case of the dual-route model of reading comprehension.

“Lexical” and “Nonlexical” Reading Routes

This use of the terms “lexical” and “nonlexical” for referring to the two reading routes seems to have originated with Coltheart (1980). Reading via the lexical route involves looking up a word in a mental lexicon containing knowledge about the spellings and pronunciations of letter strings that are real words (and so are present in the lexicon); reading via the nonlexical route makes no reference to this lexicon, but instead involves making use of rules relating segments of orthography to segments of phonology. The quotation from de Saussure with which this chapter began suggested that the orthographic segments used by the nonlexical route are single letters, but, as discussed by Coltheart (1978), that cannot be right, since in most alphabetically written languages single phonemes are frequently represented by sequences of letters rather than single letters. Coltheart (1978) used the term “grapheme” to refer to any letter or letter sequence that represents a single phoneme, so that TH and IGH are the two graphemes of the two-phoneme word THIGH. He suggested that the rules used by the nonlexical reading route are, specifically, grapheme–phoneme correspondence rules such as TH → /θ/ and IGH → /ai/.

Phenomena Explained via the Dual-Route Model

This model was meant to explain data not only from normal reading, but also facts about disorders of reading, both acquired and developmental.

Reaction times in reading-aloud experiments are longer for irregular words than regular words, and the dual-route model attributed this to that fact that the two routes generate conflicting information at the phoneme level when a word is irregular, but not when a word is regular: resolution of that conflict takes time, and that is responsible for the regularity effect in speeded reading aloud. Frequency effects on reading aloud were explained by proposing that access to entries for high-frequency words in the mental lexicon was faster than access for low-frequency words. From that it follows, according to the dual-route model, that low-frequency words will show a larger regularity effect, since lexical processing will be relatively slow for such words and there will be more time for the conflicting information from the nonlexical route to affect reading; and this interaction of frequency with regularity was observed.

Suppose brain damage in a previously literate person selectively impaired the operation of the lexical route for reading aloud while leaving the nonlexical route intact. What would such a person's reading be like? Well, nonwords and regular words would still be read with normal accuracy because the nonlexical route can do this job; but irregular words will suffer, because for correct reading they require the lexical route. If it fails with an irregular word, then the response will just come from the nonlexical route, and so will be wrong: *island* will be read as "iz-land," *yacht* to rhyme with "matched," and *have* to rhyme with "cave." Exactly this pattern is seen in some people whose reading has been impaired by brain damage; it is called surface dyslexia, and two particularly clear cases are those reported by McCarthy and Warrington (1986) and Behrmann and Bub (1992). The occurrence of surface dyslexia is good evidence that the reading system contains lexical and nonlexical routes for reading aloud, since this reading disorder is exactly what would be expected if the lexical route is damaged and the nonlexical route is spared.

Suppose instead that brain damage in a previously literate person selectively impaired the operation of the nonlexical route for reading aloud while leaving the lexical route intact. What would such a person's reading be like? Well, irregular words and regular words would still be read with normal accuracy because the lexical route can do this job; but nonwords will suffer, because for correct reading they require the nonlexical route. Exactly this pattern – good reading of words with poor reading of nonwords – is seen in some people whose reading has been impaired by brain damage; it is called phonological dyslexia (see Coltheart, 1996, for a review of such studies). This too is good evidence for a dual-route conception of the reading system.

The reading disorders just discussed are called acquired dyslexias because they are acquired as a result of brain damage in people who were previously literate. The term "developmental dyslexia," in contrast, refers to people who have had difficulty in learning to read in the first place, and have never attained a normal level of reading skill. Just as brain damage can selectively affect the lexical or the nonlexical reading route, perhaps also learning these two routes is subject to such selective influence. This is so. There are children who are very poor for their age at reading irregular words but normal for their

age at reading regular words (e.g., Castles & Coltheart, 1996); this is developmental surface dyslexia. And there are children who are very poor for their age at reading non-words but normal for their age at reading regular words and irregular words (e.g., Stothard, Snowling, & Hulme, 1996); this is developmental phonological dyslexia. Since it appears that difficulties in learning just the lexical and or just the nonlexical route can be observed, these different patterns of developmental dyslexia are also good evidence for the dual-route model of reading.

Computational Modeling of Reading

We have seen that the dual-route conception, applied both to reading aloud and to reading comprehension, was well established by the mid-1970s. A major next step in the study of reading was computational modeling.

A computational model of some form of cognitive processing is a computer program which not only executes that particular form of processing, but does so in a way that the modeler believes to be also the way in which human beings perform the cognitive task in question. Various virtues of computational modeling are generally acknowledged – for example, it allows the theorist to discover parts of a theory that are not explicit enough; inexplicit parts of a theory cannot be translated into computer instructions. Once that problem is solved and a program that can actually be executed has been written, the modeler can then determine how closely the behavior of the model corresponds to the behavior of humans. Do all the variables that influence the behavior of humans as they perform the relevant cognitive task also affect the behavior of the program, and in the same way? And do all the variables that influence the behavior of the program as it performs the relevant cognitive task also affect the behavior of humans, and in the same way? Provided that the answer to both questions is yes, studying the behavior of the computational model has demonstrated that the theory from which the model was generated is sufficient to explain what is so far known about how humans perform in the relevant cognitive domain. That does not mean that there could not be a different theory from which a different computational model could be generated which performed just as well. If that happens, the time has come for working out experiments about which the theories make different predictions – that is, whose outcomes in simulations by the two computational models are in conflict.

Of all cognitive domains, reading is the one in which computational modeling has been most intensively employed. This began with the interactive activation and competition (IAC) model of McClelland and Rumelhart (1981) and Rumelhart and McClelland (1982). This was a model just of visual word recognition, not concerned with semantics or phonology. The latter domains were introduced in the much more extensive computational model developed in a seminal paper by Seidenberg and McClelland (1989). One influence their paper had was to prompt the development of a computational version of the dual-route model: the DRC (“dual-route cascaded”) model (Coltheart et al., 1993; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001).

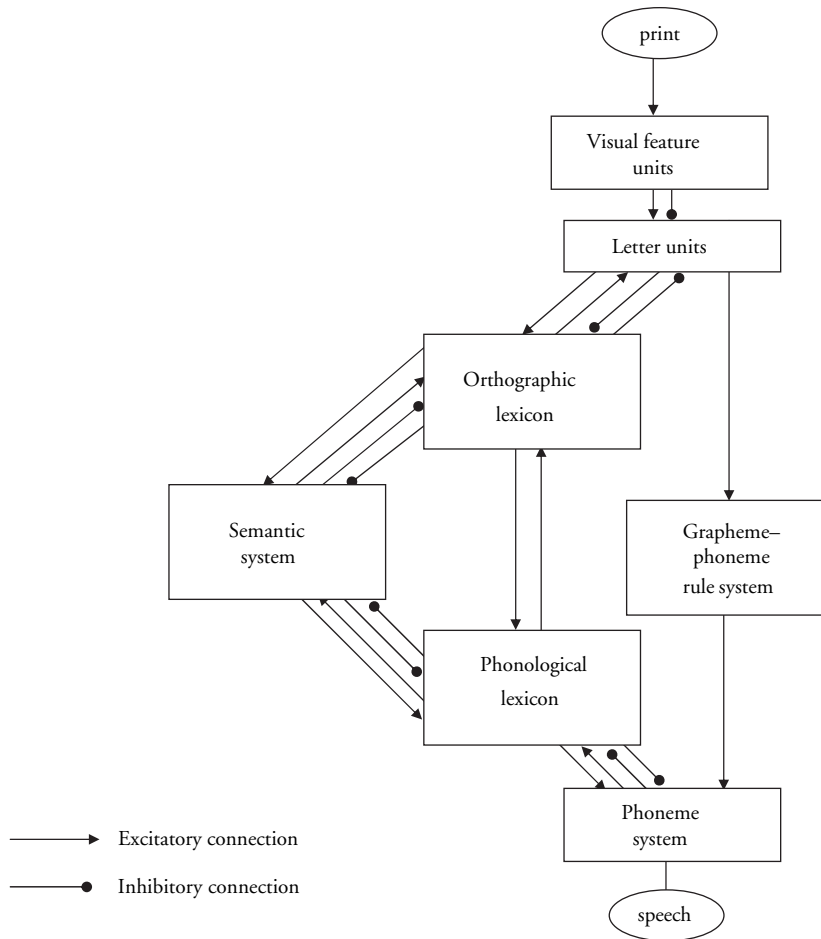
12 *Max Coltheart*

Figure 1.2 The DRC model.

The Dual-Route Cascaded (DRC) Model

The DRC is a computational model that computes pronunciation from print via two procedures, a lexical procedure and a nonlexical procedure (see figure 1.2).

The lexical procedure involves accessing a representation in the model's orthographic lexicon of real words and from there activating the word's node in the model's phonological lexicon of real words, which in turn activates the word's phonemes at the phoneme level of the model. Nonwords cannot be correctly read by this procedure since they are not present in these lexicons, but that does not mean that the lexical route will simply not produce any phonological output when the input is a nonword. A nonword such as SARE can produce some activation of entries in the orthographic lexicon for words visually similar to it, such as CARE, SORE, or SANE; this in turn can activate the phono-

logical lexicon and hence the phoneme level. Such lexically generated activation cannot produce the correct pronunciation for a nonword, but there is evidence that it does influence the reading aloud of nonwords. For example, a nonword like SARE which is similar to many entries in the orthographic lexicon will be read aloud with a shorter reaction time (RT) than a nonword like ZUCE which is similar to few (McCann & Besner, 1987).

The nonlexical procedure of the DRC model applies grapheme–phoneme correspondence rules to the input string to convert letters to phonemes. It does so in serial left-to-right fashion, initially considering just the first letter in the string, then the first two letters, then the first three letters, and so on, until it gets past the last letter in the input. It correctly converts nonwords from print to sound, and also regular words (those that obey its grapheme–phoneme correspondence rules). Irregular (exception) words are “regularized” by the nonlexical procedure – that is, their rule-based pronunciations, which will be incorrect.

Processing along the lexical route occurs as follows:

Cycle 0: set all the units for visual features that are actually present in the input string to 1; set all others to zero.

Cycle 1: every visual feature set to 1 contributes activation to all the letters in the letter units to which it is connected. The connections are inhibitory when the letter does not contain that feature, and so the activation contributed is negative; the connections are excitatory when the letter does contain that feature, and so the activation contributed is positive.

Cycle 2: what happens on Cycle 1 again happens here. In addition, every letter unit contributes activation to all the word units in the orthographic lexicon to which it is connected. The connections are inhibitory when the word does not contain that letter, and so the activation contributed from letter unit to word unit is negative; the connections are excitatory when the word does contain that letter, and so the activation contributed from letter unit to word unit is positive.

Cycle 3: everything that happens on Cycle 1 and Cycle 2 happens again here. In addition:

- (a) Feedforward: each unit in the orthographic lexicon contributes activation to its corresponding unit in the phonological lexicon.
- (b) Feedback: every word unit in the orthographic lexicon unit contributes activation back to all the letter units to which it is connected. The connections are inhibitory when the word does not contain that letter, and so the activation contributed from word unit to letter unit is negative; the connections are excitatory when the word does contain that letter, and so the activation contributed from word unit to letter unit is positive.

Cycle 4: everything that happens on Cycles 1, 2, and 3 happens again here. In addition:

- (a) Feedforward: every unit in the phonological lexicon contributes activation to all the phoneme units to which it is connected. The connections are inhibitory when the

14 *Max Coltheart*

word's pronunciation does not contain that phoneme, and so the activation contributed from word unit to phoneme unit is negative; the connections are excitatory when the word's pronunciation does contain that phoneme, and so the activation contributed from word unit to phoneme unit is positive.

- (b) Feedback: every unit in the phonological lexicon contributes feedback activation to its corresponding unit in the orthographic lexicon.

Cycle 5: everything that happens on Cycles 1, 2, 3, and 4 happens again here. In addition: every phoneme unit contributes activation back to all the word units in the phonological lexicon to which it is connected. The connections are inhibitory when the word does not contain that phoneme, and so the activation contributed from phoneme unit to word unit is negative; the connections are excitatory when the word does contain that phoneme, and so the activation contributed from phoneme unit to word unit is positive.

And so it goes. As processing cycles progress, inhibitory and excitatory influences continue to flow upwards and downwards in the way described above until the reading-aloud response is ready. How is this readiness determined? As follows. In the description of processing cycles given above, the first cycle on which the phoneme system receives any activation is Cycle 4. At the end of cycle 4, some phoneme units will be activated, but extremely weakly. As processing continues, activation of some of the phoneme units will slowly rise. Quite often, early in processing, some of the phoneme units activated will be incorrect ones. But over time as phoneme activations continue to rise it is the correct phonemes that are the most activated. A reading response is considered to be ready when phonemes have reached a critical level of activation (set to .43 when the model is being used for simulating human reading aloud). The pronunciation generated by the model is taken to consist of the most highly activated phoneme within each of the eight sets of phoneme units (one set per position) that comprise the phoneme system. The processing cycle on which that state of affairs occurs is the DRC model's reading-aloud latency for the particular letter string that was input.

Processing along the nonlexical route does not begin to operate until cycle 10. Without this time lapse after the lexical route begins to operate, the model would have serious difficulty in reading aloud irregular words. When cycle 10 is reached, the nonlexical route translates the first letter of the string into its phoneme using the appropriate grapheme-phoneme rule, and contributes activation to the phoneme's unit in the phoneme system. This continues to occur for the next 16 processing cycles. The grapheme-phoneme conversion (GPC) system operates from left to right, so eventually will move on to consider the second letter in the string as well as the first. Every 17 cycles, the GPC system moves on to consider the next letter, translate it to a phoneme, and activate that phoneme in the phoneme system. So with the letter string DESK, the GPC system has no input until cycle 10, deals with just D until cycle 27, deals with just DE from cycle 28 to cycle 44, then DES until cycle 60, DESK until cycle 76 and so on.

Computations on the lexical and nonlexical route occur simultaneously – that is, information from the visual feature level is thought of as flowing simultaneously through the lexical and the nonlexical routes and converging on the phoneme system from these two sources. Whenever the input is an irregular word or a nonword, the two sources of activation conflict at the phoneme level. If the system is to produce correct pronunciations

for irregular words and for nonwords, it will have to have a way of resolving these conflicts in favor of the correct pronunciation. Nevertheless, the model reads aloud irregular words and nonwords with high accuracy, so these conflicts are almost always resolved in a way that results in a correct pronunciation (via the interplay of inhibition and activation at various levels of the model). This depends on a judicious choice of values for the parameters of the model, such as the strengths of the inhibitory and the facilitatory connections between components of the model. If the lexical route is too strong relative to the nonlexical route, all words will be read correctly but there will be nonword reading errors. If the lexical route is too weak relative to the nonlexical route, all regular words and nonwords will be read correctly but there will be errors in reading irregular words. A delicate balance between the strengths of the two routes is needed if the model is to perform well with both nonwords and irregular words.

What the DRC Model Can Explain

One way in which Coltheart et al. (2001) evaluated the DRC model was to compare its reaction times to particular sets of stimuli to the reaction times of human readers when they are reading aloud the same stimuli. Do variables that affect human reading-aloud reaction times also affect DRC's reading-aloud reaction times? Many examples where this was so were reported by Coltheart et al. (2001). For both human readers and the DRC model:

- (a) High-frequency words are read aloud faster than low-frequency words.
- (b) Words are read aloud faster than nonwords.
- (c) Regular words are read aloud faster than irregular words.
- (d) The size of this regularity advantage is larger for low-frequency words than for high-frequency words.
- (e) The later in an irregular word its irregular grapheme–phoneme correspondence is, the less the cost incurred by its irregularity. So CHEF (position 1 irregularity) is worse than SHOE (position 2 irregularity), which is worse than CROW (position 3 irregularity).
- (f) Pseudohomophones (nonwords that are pronounced exactly like real English words, such as brane) are read aloud faster than non-pseudohomophonic nonwords (such as brene).
- (g) Pseudohomophones derived from high-frequency words (e.g., hazz) are read aloud faster than pseudohomophones derived from low-frequency words (e.g., glew).
- (h) The number of orthographic neighbors a non-pseudohomophonic nonword has (i.e., the number of words that differ from it by just one letter), the faster it is read aloud.
- (i) The number of orthographic neighbors a pseudohomophone has does not influence how fast it is read aloud.
- (j) The more letters in a nonword there are the slower it is read aloud; but number of letters has little or no effect on reading aloud for real words.

16 *Max Coltheart*

The DRC model was also used to simulate acquired dyslexias. Surface dyslexia was simulated by slowing down rate of access to the orthographic lexicon: this lesioned DRC made regularization errors with irregular words, more so when they were low in frequency, just as is seen in surface dyslexia, whereas its reading aloud of regular words and nonwords remained normal, as in the pure cases of surface dyslexia (Behrmann & Bub, 1992; McCarthy & Warrington, 1986). Phonological dyslexia was simulated by slowing down the operation of the nonlexical route: this lesioned DRC still read words correctly, but misread nonwords, especially if they were nonpseudohomophones, as in the case of phonological dyslexia.

Thus, the DRC model can explain an impressively large number of findings from studies of normal and disordered reading, far more than any other computational model of reading. Nevertheless, Coltheart et al. (2001) drew attention to a number of limitations of the current implementation of the DRC model: its procedure for performing the lexical decision task was crude, it was not applicable to the pronunciation of polysyllabic words or nonwords, it did not offer any account of one popular paradigm for studying reading (masked priming), the difference between word and nonword reading RTs by the model was probably implausibly large, the amount of variance of word reading RTs that the model could account for, though always significant, was disappointingly low, and the implemented model has nothing to say about semantics. A new version of the DRC model that will correct these and other shortcomings of the existing model is under development.

Connectionist and nonconnectionist modeling

This chapter distinguishes between connectionist models of reading (such as the models of Seidenberg & McClelland, 1989, and Plaut, McClelland, Seidenberg, & Patterson, 1996) and nonconnectionist models of reading (such as the DRC model). The description of the DRC model in Coltheart et al. (2001) uses the term “connection” and the model in fact “contains” about 4.5 million connections, in the sense of the term “connection” used by Coltheart et al. (2001). However, in the DRC model, connections are just expository devices used for talking about how the modules of the model communicate with each other. One could expound this in other ways without using the term “connection.” In contrast, in connectionist models, the connections are often thought of as neuron-like, the models are referred to as neural networks, and terms like “biologically inspired” or “neurally plausible” are often applied. Here a connection is something that is physically realizable as an individual object, in contrast to the DRC model in which there is no such sense to the term.

A second major difference between connectionist and nonconnectionist modeling, at least as those trades have been practiced up until now, is that connectionist models have typically been developed by applying a neural-net learning algorithm to a training set of stimuli, whereas the architectures of nonconnectionist models have typically been specified by the modeler on the basis of the empirical effects that the model is meant to explain.

The Seidenberg and McClelland (1989) connectionist computational model of reading is often presented as an alternative to the dual-route model. Indeed, claims such as “The dual-route model has been more recently questioned by a plethora of single-route computational models based on connectionist principles” (Damper & Marchand, 2000,

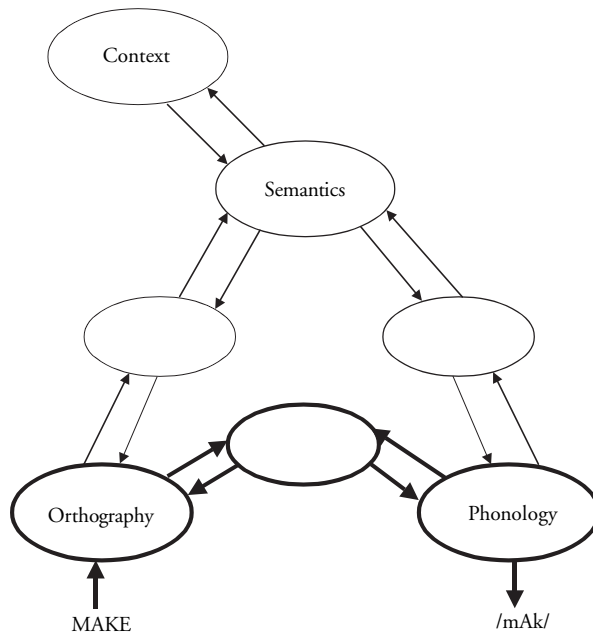


Figure 1.3 The Seidenberg and McClelland (1989) model. (The implemented model is in bold-face type.)

p. 13) are common in the literature. But that was not the view of the authors themselves. They were clear about this: “Ours is a dual route model,” they stated (Seidenberg & McClelland, 1989, p. 559).

This is perfectly evident from their diagram of their model (Seidenberg & McClelland, 1989, figure 1, reproduced as figure 1.3 here): it explicitly represents two distinct routes from orthography to phonology, one direct and the other via meaning, and explicitly represents two distinct routes from orthography to semantics, one direct and the other via phonology. One of the two routes for reading aloud (the one via semantics) can only be used for reading words aloud; it would fail for nonwords. The other (nonsemantic) route for reading aloud is required if the stimulus is a nonword. This model has come to be called the triangle model, perhaps because of the reference in Seidenberg and McClelland (1989, p. 559) to “the third side of the triangle in Figure 1.” More than one subsequent model has been referred to as the triangle model despite being different from Seidenberg and McClelland’s model. So far there have been seven different triangle models, an issue discussed later in this chapter.

What is it that has led to this widespread misunderstanding? The answer is clear: a failure to distinguish between the following two claims:

- (a) It is possible for a single processing system to correctly read aloud all irregular words and all nonwords.
- (b) The human reading system possesses only one procedure for computing pronunciation from print.

Seidenberg and McClelland (1989) did make claim (a). But they did not make claim (b); indeed, as the quotation in the previous paragraph indicates, they repudiated claim (b). That is why theirs is a dual-route model of reading aloud.

This seminal model turned out not to be able to offer a good account of how people read nonwords aloud because its accuracy on this task was far less than the accuracy that human readers show (Besner, Twilley, McCann, & Seergobin, 1990). The suggestion (Seidenberg & McClelland, 1990, p. 448) that this was because the database of words on which the model was trained was too limited and did not contain enough information for nonword reading to be learned from it was shown to be incorrect by Coltheart et al. (1993). They developed a GPC rule-learning algorithm and applied it to the Seidenberg–McClelland training set. The rule set that this algorithm learned from that training set was then used with 133 nonwords from Glushko (1979). Whereas the Seidenberg and McClelland model scored only 68% correct on a subset of 52 of these nonwords, the DRC read 97.9% of these correctly. This shows that the information needed to learn to be an excellent nonword reader is actually present in the model's database, and so "the poor performance of the PDP model in reading nonwords is a defect not of the database but of the model itself" (Coltheart et al., 1993, p. 594). Hence, as noted by Plaut (1997, p. 769) and (Plaut et al., 1996, p. 63), the Seidenberg and McClelland model did not succeed in providing evidence that it is possible for a single processing system to correctly read aloud all irregular words and all nonwords.

Nevertheless, it might well be possible to devise a single processing procedure that can correctly read aloud all irregular words and all nonwords. Plaut et al. (1996) sought to devise such a procedure via training a connectionist network similar in overall architecture to that of the network of Seidenberg and McClelland shown in figure 1.3 (it was, for example, a dual-route model in just the same sense that Seidenberg and McClelland viewed their model as a dual-route model, though training was carried out on only one of the two routes), but differing from the Seidenberg and McClelland model in a number of ways, including in the forms of orthographic and phonological representations used in the network. Input units, which were distributed representations in the Seidenberg and McClelland model, became local representations (each representing a grapheme). Output units, which were distributed representations in the Seidenberg and McClelland model, became local representations (each representing a phoneme).

Plaut et al. (1996) actually presented three different though related models – that is, a second, third and fourth triangle model, the first triangle model being that of Seidenberg and McClelland (1989):

Model 1: purely feedforward, 105 grapheme units, 100 hidden units, 61 phoneme units.

Model 2: as for Model 1 but with feedback from phoneme units back to hidden units: an attractor network.

Model 3: as for Model 1 but adding (unimplemented) external input to the output units, so as to mimic what could happen if there were an implemented semantic system activated by orthography and in turn activating phonology. This approach, discussed further below, was pursued in an attempt to simulate acquired surface dyslexia.

How well do these models read nonwords? Model 1 (which after training scored 100% on reading the 2,972 nonhomographic words in the training set) did quite well on nonword reading (see table 3 of Plaut et al., 1996), almost as well as human readers. However it still fails with items like JINJE, the reason being that there is no word in the training corpus that ends with the final grapheme of this nonword. It follows that careful selection of nonwords which exploits such gaps in the training corpus would produce a set of nonwords on which the model would score at or close to zero. Human readers would be vastly superior to the model on such nonwords. Results with nonword reading by Model 2 were similar, though its nonword reading was slightly worse than that of Model 1. The JINJE problem remained.

Given this work by Plaut et al. (1996), what are we to say about the two claims mentioned above? These claims were:

- (a) It is possible for a single processing system to correctly read aloud all irregular words and all nonwords.
- (b) The human reading system possesses just one procedure for computing pronunciation from print.

Although nonword reading was better by the PMSP models than by the SM model, the PMSP models still do not read nonwords correctly in the sense of “as well as human readers do,” since it is not difficult to devise nonwords that human readers read well and the PMSP models read wrongly: there is no sense in which reading JINJE to rhyme with “wine” (as the PMSP models do) could be regarded as correct. So claim (a) remains without support. And no current model of reading aloud makes claim (b). Hence at present it is reasonable to regard both claims as false.

However, the work on simulation of surface dyslexia using Model 3 has an interesting implication for these claims. Indeed, in general simulation of disordered rather than normal reading it has been particularly crucial in recent years for comparative evaluation of computational models of reading. Hence much of the following discussion of dual-route modeling will focus on the application of such models to the explanation of disordered reading.

Simulating disordered reading with the triangle models

Simulating acquired surface dyslexia. Acquired surface dyslexia (Marshall & Newcombe, 1973; Patterson, Marshall, & Coltheart, 1985) is a reading disorder, caused by brain damage, in which there is selective impairment of the ability to read irregular words aloud with relative sparing of regular word and nonword reading. Many cases are not normal at regular word and nonword reading; I will focus here, as did Plaut et al. (1996), on two particularly pure cases, KT (McCarthy & Warrington, 1986) and MP (Behrmann & Bub, 1992). Both showed virtually normal accuracy in reading aloud regular words and nonwords, but were impaired at reading irregular words, especially when these were low in frequency (KT: high frequency 47%; low frequency 26%; MP: high frequency 93%; low frequency 73%).

20 *Max Coltheart*

Computational models are meant to be able to explain impaired reading as well as normal reading: that is, it should be possible to artificially lesion these models so that their patterns of preserved and impaired reading correctly match such patterns seen in various forms of acquired dyslexia. Plaut and colleagues therefore investigated whether there was any way of lesioning any of their three models that would lead to impaired irregular word reading with preserved regular word and nonword reading.

This was investigated by studying the effects of deleting various proportions of the connections in the implemented orthography-to-phonology pathway, or various proportions of the hidden units, in Model 2. This was not successful in simulating the more severe patient KT: any lesion that produced accuracies of around 26% for low-frequency irregular words also produced very poor performance with nonwords, whereas KT was perfect at reading nonwords. It was therefore not possible to simulate acquired surface dyslexia just with the implemented part of the model.

So Plaut et al. turned from Model 2 to Model 3, which has an unimplemented component (semantic input to the phonological output level). With sufficient training, Model 3 does well with irregular words, regular words, and nonwords. What is crucial here, though, is the competence of the implemented (orthography-to-phonology) part of Model 3. When it is trained without semantics (this is Model 1), it learns to read irregular words perfectly and nonwords very well. But this is not the case when it is trained with concurrent semantic input. Low-frequency irregular words are never learned perfectly by the direct orthography-to-phonology pathway here: for this pathway operating on its own, accuracy for low-frequency irregular words is about 70% after 400 epochs of training and then declines down to about 30% correct after 2,000 epochs. Performance with high-frequency irregular words is almost perfect at 400 epochs, but further training progressively worsens performance with these words, down to about 55% at epoch 2,000. Regular word and nonword performance is almost perfect at epoch 400 and remains at that level with further training to epoch 2,000.

If training is stopped at 400 epochs, and semantic input to the system is then deleted, performance is good with regular words, nonwords, and high-frequency irregular words, but somewhat impaired with low-frequency irregular words; that matches the surface dyslexic pattern shown by MP.

If training is stopped at 2,000 epochs, and semantic input to the system is then deleted, performance is good with regular words, and nonwords, impaired with high-frequency irregular words, and very poor with low-frequency irregular words; that matches the surface dyslexic pattern shown by KT.

The suggestion here is that the cause of acquired surface dyslexia is semantic damage, and that the more the patient had relied on semantic input for reading aloud pre-morbidly, the more severe the surface dyslexia will be when semantic damage occurs. The implication is that, even if it is possible for a single processing system to correctly read aloud all irregular words and all nonwords, most human readers do not possess such a system.

Because there are patients with severe semantic damage who can read irregular words with normal accuracy (e.g., Cipolotti & Warrington, 1995; Lambon Ralph, Ellis, & Franklin, 1995; Schwartz, Saffran, & Marin, 1980a; see also Gerhand, 2001), Plaut

et al. (1996, p. 99) had to suppose that some people learn to read without any support from semantics and so can read all irregular words without recourse to semantics. But in other work using the triangle models this supposition has been abandoned:

It is important to note that, because this version of the triangle model assumes a causal relationship between semantic impairment and surface dyslexia, its adequacy is challenged by any observations of semantically impaired patients whose reading does not reveal a surface dyslexic pattern. (Fushimi et al., 2003, p. 1656)

A degraded semantic system will inevitably impair the ability to “know” a letter string . . . as belonging to the repertoire of real words. (Rogers, Lambon Ralph, Hodges, & Patterson, 2004, p. 347)

According to Model 3 as it is applied to the analysis of surface dyslexia, intact human readers possess two routes from print to speech. Let’s call these, theory-neutrally, Route A and Route B. Properties of these routes are:

- (a) Route A can correctly read aloud all known words (regular or irregular) but cannot read nonwords aloud correctly.
- (b) Route B can correctly read aloud all regular words and all nonwords, but will misread X% of irregular words.

This connectionist dual-route model of reading aloud differs from the nonconnectionist dual-route DRC model of reading aloud (Coltheart et al., 2001, discussed below) only with respect to the value of X. According to Plaut et al. (1996), premorbidly X can on rare occasions be zero (the patients referred to above who are normal at irregular word reading but have severe semantic impairments) but typically is not and can be at least as high as 64% (patient KT’s overall error rate on irregular words). According to the DRC model, X is always 100%.

So, while it is of course logically possible that the system humans use for reading aloud has a single-route architecture, there are no theoretical proposals embodying such an architecture that can escape refutation from available data from studies of normal and impaired readers. All the models are dual-route models. Current and future theorizing is and will be about the details of what these two routes are actually like.

Simulating acquired phonological dyslexia

Harm and Seidenberg (2001) used another connectionist triangle model in work attempting to simulate acquired phonological dyslexia. In their view, this form of acquired dyslexia is always caused by a phonological impairment. Therefore, after training their model until it was performing well in reading words and nonwords, they lesioned the phonological component of the model by adding random noise each time the units in that component were being updated. This harmed nonword reading more than word reading and so simulated phonological dyslexia. However, this explanation of acquired

22 *Max Coltheart*

phonological dyslexia predicts that cases of acquired phonological dyslexia without the presence of a phonological impairment will not be seen, and this prediction is incorrect. Dérouesné and Beauvois (1985), Bisiacchi, Cipolotti, and Denes (1989), and Caccappolo-van Vliet, Miozzo, & Stern (2004) have all reported cases of acquired phonological dyslexia with preserved phonological processing.

As we have seen, the development of connectionist triangle models of reading has been considerably influenced by attempts to simulate acquired dyslexia; and this approach has also been applied to the simulation of developmental dyslexia.

Simulating developmental dyslexia. Harm and Seidenberg (1999) developed a model in which to simulate developmental reading disorders. Their particular triangle model differed from all earlier triangle models in a number of ways:

- (a) Learning in the phonological units was assisted by the presence of a set of cleanup units attached to the phonological units.
- (b) The phonological units represented phonetic features, not phonemes.
- (c) The orthographic units represented letters, not graphemes.
- (d) Positional coding of orthography was relative to the vowel in the input string, rather than absolute.

After training, the model achieved satisfactory levels of performance in reading the irregular words in the training set, and also in reading nonwords (though again performance seemed slightly inferior to human nonword reading).

Harm and Seidenberg (1999) were specifically interested in attempting to simulate developmental dyslexia. Having shown that their triangle model was capable of learning to read adequately, they then investigated ways of impeding its learning that might result in either of two different subtypes of developmental dyslexia, one in which nonword reading is selectively affected (developmental phonological dyslexia) and another in which irregular word reading is selectively affected (developmental surface dyslexia; Harm and Seidenberg preferred the term “reading delay dyslexia” because they believed that the reading of children with developmental surface dyslexia is just like the reading of younger children who are learning to read normally).

Because Harm and Seidenberg (1999) believed that developmental phonological dyslexia is always caused by the child having a phonological processing deficit, their approach to simulating developmental phonological dyslexia involved lesioning their model’s phonological system. This was done in two different ways:

- (a) Mild phonological impairment: a slight degree of weight decay was imposed on the phonetic feature units throughout training.
- (b) Moderate phonological impairment: in addition to the weight decay, the cleanup units were removed from the network, as were a random 50% of the interconnections between the phonetic feature units.

Both types of lesioning did impair the model’s ability to learn to read nonwords. But when this impairment was more than mild, the ability of the model to learn to read words

was also impaired. Hence what could not be simulated here was pure severe developmental phonological dyslexia (where “pure” means that word reading is in the normal range and “severe” means the impairment of nonword reading was more than mild). That raises the question: does one ever see pure severe developmental phonological dyslexia in human readers? A number of such cases have been reported (see e.g. Campbell & Butterworth, 1985; Funnell & Davison, 1989; Holmes & Standish, 1996; Howard & Best, 1996; Stothard et al., 1996). Hence these data from developmental cognitive neuropsychology provide a challenge for the Harm and Seidenberg (1999) connectionist model of reading.

Developmental surface dyslexia (“reading delay dyslexia”) was simulated in the work of Harm and Seidenberg (1999) by reducing the number of hidden units in the network from 100 to 20, and also by reducing the network’s learning rate. Both types of developmental damage to the network harmed the learning of irregular words more than the learning of nonwords; but in both cases the learning of nonwords suffered too. Thus it was not possible to simulate “pure” developmental surface dyslexia (i.e., impaired irregular word reading with *normal* nonword reading). However, pure developmental surface dyslexia is seen in human readers (Castles & Coltheart, 1996; Hanley & Gard, 1995; Goulandris & Snowling, 1991). Hence again these data from developmental cognitive neuropsychology do not provide support for the Harm & Seidenberg (1999) connectionist model of reading.

Conclusions

Reading theorists have reached unanimity concerning the existence in the human reading system of two separate procedures for reading aloud – that is, dual routes from print to speech. One of these processing routes is usable only when the stimulus to be read is a real word; it cannot read nonwords. The other route can read all nonwords and regular words; there is still some dispute concerning how well it reads irregular words.

These dual-route models differ in terms of whether they are connectionist models such as the triangle models or nonconnectionist models such as the DRC model. At present the data favor the nonconnectionist approach. The DRC model does a good job of simulating patterns of acquired dyslexia, which the connectionist models have not succeeded in doing. Nor have the connectionist models succeeded in accounting for developmental reading disorders, whereas the DRC model is compatible with everything we currently know about these disorders. Finally, none of the connectionist models can explain all of the phenomena from studies of normal reading listed above (see the section “What the DRC Model Can Explain”), whereas all of these can be simulated by the DRC model.