

The Psychology of Reading

On first appearances, the task of reading sentences like this one is seemingly so simple that it requires no explanation: The reader – in this instance, you – simply looks at the words on the page, identifying them in turn, and then (somehow) combining their meanings to understand the contents of each new sentence. But this subjective experience is misleading. The science of reading has shown that reading is one of – if not *the* – most difficult activities in which we routinely engage but for which we have no biological predisposition. Reading thus stands in contrast to spoken language, which is another difficult activity that we routinely engage in but for which we are genetically predisposed. It is an activity that arguably requires a significant portion of the brain and the cognitive systems that it supports to execute, including the systems that support vision and attention (to identify the printed words), long-term memory (the repository of our world knowledge), language (to construct the meanings of sentences), and even motor control (to move the eyes from one word to the next). The main goal of the *psychology of reading*, therefore, is to develop an account of how we, as readers, can coordinate these activities to convert the sequences of ink marks on a printed page into the potentially infinite number of ideas that can be conveyed through writing.

Although the “psychology of reading” might sound strange to the uninitiated (given that many people view psychology as the study and treatment of mental health problems), it refers to a subdiscipline of a particular branch of psychology called *cognitive psychology*. As its name might suggest, cognitive psychology is the study of the basic perceptual and mental processes (i.e., cognition) that make us who we are (for an introductory text, see Eysenck & Keane, 2015). A short list of these processes include vision, attention, memory, language, reasoning, problem solving, consciousness, and any of the many tasks that we humans engage in, including reading and many others (e.g., driving automobiles, solving syllogisms, playing chess, etc.).

Because of the sheer complexity of what needs to be explained by the psychology of reading, researchers in the field have developed a variety of ingenious methods to study the behaviors of reading, the mental processes and representations that are produced by or are the product of those behaviors, and the neural systems that support those mental processes and behaviors. These research methods are the building blocks of subsequent chapters in this book, and as such, they will be discussed in some detail below. However, an equally important tool for reading researchers are the theories that guide the experimental research. Because the psychology of reading is an advanced science, the field has benefited from the development of several formal *models*, or theories that have been implemented using mathematical equations and computer programs, that both describe and simulate the various processes that are involved in reading (e.g., the identification of printed words), often with remarkable accuracy (for an introduction and review of these models, see Reichle, 2021). Because these reading models play a central role in motivating new empirical research, this chapter will also provide a brief introduction to two prominent examples and briefly discuss how models are compared and evaluated.

Before we review the methods and models of reading research, however, it is first necessary to have a clear understanding about what reading is, and how it is related to both spoken language and other forms of communication. The next section of this chapter does exactly that. But before we launch into this discussion, we also want to say briefly what this book is about, and where subsequent chapters will take us. As the title of this book suggests, it is about the psychology of reading, and more specifically, what has recently been learned about the psychology of reading from experiments on and models of the reading of one language and writing system – Chinese.

At this point, a perfectly natural question to ask might be: Why Chinese? The short answer to this question is that most research that has been done to understand the psychology of reading has focused on how people who speak European languages go about reading alphabetic texts (e.g., English, Spanish, Russian, etc.). Although this research has been highly informative and has taught us a considerable amount about reading (e.g., see the monographs by Crowder & Wagner, 1992; Dehaene, 2009; Perfetti, 2005; Rayner & Pollatsek, 1989; Rayner et al., 2012; Reichle, 2021; Seidenberg, 2017; Taft, 1991; Wolf, 2008; see also the edited volumes by Coltheart, 1987; Klein & McMullen, 2001; Pollatsek & Treiman, 2015; Snowling & Hulme, 2005), this endeavor has until recently largely avoided the possible implications of languages and writing systems that differ markedly

from the ones that have been studied – languages and writing systems like those required to read Chinese. This oversight is perhaps a bit like an ornithologist who concludes that all birds can fly because he or she has never encountered a penguin, emu, or ostrich. This oversight has at least in some instances resulted in theoretical assumptions about the psychology of reading that are likely to be incomplete or even incorrect. Subsequent chapters in this book will explore some of these (potentially) faulty assumptions, but for now, we will continue with our exposition of reading, language, and communication.

1.1 **Communication, Language, and Reading**

Although anyone reading this book will be intimately familiar with both reading and language, it is important to distinguish between the two to avoid unwarranted assumptions. To begin with, spoken language is a natural endowment of the human species – a capacity to communicate that develops in all neurologically normal children, irrespective of their culture, and with only minimal explicit instruction (Deacon, 1997; Pinker, 2015). All spoken languages also share several necessary and sufficient features that are not shared by other forms of animal communication.

The first of these features is that language is *generative* (Chomsky, 1959). By this, we simply mean that, even if one only knows a finite number of words, it is still possible to express an infinite number of ideas. This unlimited power of expression is possible because *grammar*, or the rules that permit individual words to be combined into the larger units of meaning corresponding to phrases and sentences, allow for the expression of an unlimited number of novel ideas – ideas that vary along the continuum from being simple to tremendously complex, or that differ from other ideas in often subtle ways. And if one ignores the many non-linguistic restrictions on language comprehension (e.g., one's attention span or the degradation in understanding caused by background noise), then these sentences can in principle be of unlimited length. (A good example of this is Mike McCormack's single-sentence, 232-page novel, *Solar Bones*, 2016, which won the 2018 International Dublin Literary Award.)

A second feature of human language that also contributes to its generativity is the fact that words are *symbolic* in nature, allowing an infinite number of objects and concepts to be referenced by arbitrary combinations of sounds (i.e., *phonemes*) and visual forms (i.e., *graphemes*) within a given language. This point is immediately obvious if one compares the words for a given concept, for example “cat,” across languages; in English,

the word is written “cat” and pronounced /kæt/, whereas in Italian it is written as “gatto” and pronounced as /'gat:to/, while in Chinese it is written as “猫” and pronounced as /mao¹/.¹ These three examples illustrate how the mappings between the thing or concept being referred to, on the one hand, and the symbols that are used to represent the thing or concept, on the other, are completely arbitrary, having been established through historical accident and convention within a group of language speakers. Importantly, the symbolic nature of language adds to its generativity in that new words can and constantly are being coined. For example, consider one word that, although it is in worldwide circulation now, was largely unheard of until 2020: “Covid-19.”

The generative and symbolic nature of human language also affords its third and final feature – that the speakers of language can refer to things or situations outside of the immediate present. This allows one to speak, for example, of something that happened in the past, or that might be anticipated to happen in the future, or in the context of fantasy and science fiction stories, situations or things that are not possible in the real world. In short, languages allow their speakers to traverse time, escaping the bounds of whatever is happening in the immediate present.

These three features of human language – the fact that it is generative, symbolic, and allows one to reference the past and future – stand in stark contrast to animal communication. For example, consider the warning signals that are issued by vervet monkeys in response to different predator threats (Seyfarth et al., 1980). Although these monkeys use different calls to warn their compatriots of eagles, snakes, and leopards, these calls cannot be combined to convey more complex or novel warnings (e.g., something equivalent to “Ignore the eagle in the distance because a leopard is approaching!”). Nor can the monkeys issue warning about novel threats (e.g., “Avoid the humans!”). And finally, the calls are only issued in response to immediate danger, and cannot, for example, be issued in advance (e.g., “Beware of the leopard that will return later.”). Thus, although vervet monkeys use their warning calls to communicate important information to other vervet monkeys, this communication is extremely rigid, being fixed to a small number of messages that are directly related to the immediate present. As far as we know, these same

¹ The International Phonetic Alphabet (see Akmajian et al., 2010) will be used here and in subsequent chapters to represent all examples of phonological forms or pronunciations of words. And as will be discussed in more detail in Chapter 2, the Chinese spoken languages are tonal, with each written character having an associated tone or pitch that can be represented by a number. In the example given, the 1 represents a flat tone (i.e., one that does not change).

limitations seem to apply to all other forms of animal communication, making human language truly unique in the animal kingdom.

With this background on what language is and is not, it is now possible to contrast spoken language, on one hand, with written language (and reading), on the other.² As mentioned previously, the capacity to use spoken language is part of our genetic endowment. As such, the capacity to speak has been subject to evolutionary pressure and likely emerged over millions of years (Deacon, 1997). And during the past few millennia, spoken languages have continued to evolve among populations of speakers within different geographical regions, resulting in the roughly 6,500 languages that are spoken in the world today. The evolution of these languages can often be traced using comparative methods, allowing, for example, for a comprehensive understanding of how modern English and German diverged from a single common language to become two distinct languages.

In contrast to spoken language, the capacity to read and write are relatively recent cultural inventions in that the best available evidence suggests have only been around for approximately 5,500 years (Robinson, 1995). Like many other significant cultural inventions (e.g., legal and political systems), the insights that led to the capacity to read and write likely occurred independently across different cultures separated by vast geographic distances. Robinson provides summary evidence, for example, that economic transactions were likely being recorded on clay tablets by 3,500 BCE, with these simple notations changing to cuneiform within a few hundred years. And independently of that, other very different forms of writing seem to have developed at other locations around the world. Key milestones in the development include the emergence of hieroglyphs in ancient Egypt by 3,000 BCE, the emergence of various scripts in and around the Indus valley and Aegean sea (e.g., Linear A and B) by 2,500–1,500 BCE, the use of ideograms (the early precursors to characters) in China by 1,200 BCE, the development of the Phoenician alphabet by 1,000 BCE, the subsequent adoption and modification (e.g., addition of letters representing vowels) of the Phoenician alphabet by the Greeks around 730 BCE, and evidence of Mesoamerican scripts and hieroglyphs by 600 BCE.

Despite their tremendous variety, what all these early writing systems share with modern writing systems is their capacity to represent the sounds

² This discussion will ignore various forms of sign language that are used by the deaf. It is worth noting, however, that these languages are true languages, and that as such, what was said about spoken languages is equally applicable to sign languages. And in the same vein, although our discussion of reading will ignore the reading of braille, it is also a true writing system.

and meanings of spoken language, fixing the information in media that allows for a permanent record of not only financial transactions, but also of history and culture, as well as religious texts, poetry, and literature. The significance of this new technology cannot be overstated. As the astronomer Carl Sagan (1980) rightly noted, this capacity is “perhaps the greatest of human inventions, binding together people who never knew each other,” allowing one to be “inside of the mind of another person, maybe somebody dead for thousands of years.”

The challenge, therefore, is to develop methods for understanding this most amazing of cognitive tasks. Fortunately, this has been possible; the experimental methods and technologies of modern cognitive psychology and neuroscience have allowed reading researchers to make informed inferences about mental processes that support reading by measuring their neural correlates (e.g., patterns of brain activity) and the overt behavior required to perform various types of reading-related tasks (e.g., pronouncing words aloud).³ Section 1.2 will provide a brief introduction to the main methods and technologies that are available to do this and that have been productively used to advance the science of reading. However, please note that this introduction is not intended to be comprehensive in scope, but instead only provides the minimal background that might be required to understand the remaining chapters of this book. Readers with backgrounds in either cognitive psychology or neuroscience may therefore opt to skip ahead to those chapters.

1.2 Methods to Study Reading

The psychology of reading has a long history of scientifically rigorous experimental investigation (e.g., Huey’s, seminal *The Psychology and Pedagogy of Reading*, 1908, is perhaps the earliest comprehensive account). It is therefore not surprising that many of the experimental methods that have been used to study reading also have a long history. This section will review those methods, as well as new technologies that have only recently made it possible for researchers to study the internal behaviors of a reader’s mind or brain – the patterns of cortical activity that can be measured using the brain’s electrophysiology and metabolism.

³ We will avoid the thorny philosophical issue of specifying precisely how the mind differs – if it even does – from the brain, but will instead simply acknowledge that it is useful to think about the mind as being a more abstract description of how the brain operates. This approach is analogous to how one might think about a computer program as being an abstract description of how a computer operates (e.g., see Coltheart, 2012).

The earliest methods used to study reading involve simple behavioral tasks that, despite their simplicity, were used with ingenuity to good effect. For example, someone unfamiliar with how reading research is conducted might suggest the simple task of having participants read extended passages of text for some amount of time and then answering questions to gauge their understanding of the text's contents. This task can be used to determine the maximal reading speed or what a reader is likely to remember about a given text, but unfortunately says very little about the moment-to-moment inner workings of the mind, or how the processing of words and sentences allows a reader to construct the mental representations that are necessary to answer comprehension questions. Such insights require more sophisticated methods, typically involving tasks coupled with experimental designs that allow researchers to focus on one specific process of interest.

For example, because the key component of reading is the identification of printed words, or *lexical access*, many of the earliest techniques for studying reading were designed to shed light on this process. For example, one technique called the *perceptual-identification task* involves displaying a word under well-controlled viewing conditions using a device called a *tachistoscope*. This device contains a camera shutter that allows a word (usually printed on a card) to be displayed for a precise amount of time under specific lighting conditions. By asking participants to name or otherwise identify words that are displayed using a tachistoscope, it is possible to determine the minimum amount of time that is needed to identify a word, as well as the types of information that might be extracted when participants make errors. For example, one early and theoretically important finding from such experiments that has also withstood the scrutiny of time is the *word-superiority effect*. This effect is the seemingly paradoxical finding that a letter displayed in the context of a word can be identified more rapidly and accurately than a letter displayed in isolation (Reicher, 1969; Wheeler, 1970). For example, the letter "k" can be identified more efficiently if it is displayed in the context of "work" than if "k" is displayed in isolation. This suggests that the perception of letters is somehow facilitated or supported by the processing of the word in which they occur, and explanations of the effect typically refer to the "top-down" influence of word representations in memory influencing the visual perception of letters, with letters that are displayed in words benefiting from this additional support (e.g., see McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982).

Over the last few decades, the behavioral tasks that are used to study word identification have been refined into a small set of standardly used tasks. Although each of these tasks is useful, it is fair to say that all of these

tasks also have their limitations. For example, the first of these tasks is *naming*, wherein participants simply pronounce words that are displayed on a computer monitor as rapidly and accurately as possible, with a microphone being used to detect the onset time of the naming response (Balota et al., 2007; Schilling et al., 1998). Although this task is (at least somewhat) natural and can under some conditions guarantee that participants have accessed a word's pronunciation from memory (see, e.g., Rossmeissl & Theios, 1982), it is also clear that words can be sounded out and thus performance in this task can have little to do with lexical access. (The easiest demonstration of this is the fact that you can pronounce non-words like "fark" that, by virtue of being non-words, are not represented in the lexicon.) Another limitation of the naming task is that words beginning with certain phonemes (e.g., voiced consonants, which are more likely to trigger the voice key) are more likely to be named more rapidly and/or accurately than words beginning with other phonemes (Rastle & Davis, 2002).

A second standardly used task is *lexical decision*, wherein participants view a series of words and non-words displayed one at a time on a computer monitor, with the task of indicating as rapidly and accurately as possible via button presses whether each letter string is a word or non-word (Balota et al., 2007; Schilling et al., 1998). Because both "word" and "non-word" responses are registered using button presses, this task avoids much of the messiness of naming. But because the task requires binary decisions, the task is subject to strategies that reflect a variety of variables, including the relative proportion of words versus non-words being used in the experiment (West & Stanovich, 1982), and the degree to which the non-words resemble words (e.g., it is easier to decide that the consonant string "rxwmq" is a non-word than it is to decide that the pseudo-homophone "brane" is a non-word; Van Orden et al., 1990). For those reasons, the lexical-decision task has been criticized on the grounds that it often measures more than just lexical access (e.g., Balota & Chumbley, 1984).

Given the limitations of the naming and lexical-decision tasks, one might ask why researchers do not simply ask participants to somehow indicate the meaning of a word, as a way of guaranteeing that lexical access has occurred. One task that has been developed to do exactly that is *semantic verification*, wherein participants indicate (usually via button presses) whether or not each word in a series has some particular semantic attribute (e.g., Lewellen et al., 1993). Participants might be asked, for example, to indicate whether each of a series of words refers to something animate. For example, the sequence "cat," "hammer," and "sink" would be expected to elicit "yes," "no," and "no" responses, respectively.

Although the semantic-verification task avoids most of the pitfalls of naming and lexical decision, it also requires binary decisions that can bias responses and thus has also been criticized for being unnatural (e.g., see Van Orden, 1987).

One final task that avoids the criticism of being unnatural uses the measurement of eye movements during natural reading – an approach that is most often referred to as *eye tracking* (Rayner, 1979; for reviews, see Rayner, 1998, 2009). This technique obviously requires an *eye tracker*, or device that measures the position of a reader's eyes as they read text that is displayed on a computer monitor. The most widely used eye trackers today, for example, typically measure the position of a reader's dominant eye once every millisecond as they read sentences or passages of text, allowing the experimenter to reconstruct a variety of different measures that reflect (on average, across a sample of participants) how long and often each word is looked at during reading. Using these word-based measures, it is then possible to make informed inferences about what is happening in the mind of a reader because there is a tight coupling between the eye and mind during reading (Reingold et al., 2012). Moreover, because the task is natural (i.e., doesn't require binary decisions or other secondary task demands) and is both non-intrusive and highly sensitive (i.e., fixation durations and locations can be measured very accurately), the only real drawback of this approach is its inherent complexity. In other words, to interpret a pattern of eye movements as they move through a sentence, one must make informed inferences about how visual processing, attention, word identification, and sentence comprehension drive the moment-to-moment movement of the eyes – the inner workings of the mind that are of interest to reading researchers.

Although each of the behavioral methods differ in important ways, and although each has its merits and limitations, all have proven useful, and many of the key findings have been documented using more than one of the methods, with this convergence further validating the methods. These key findings will be discussed in subsequent chapters of this book, but for now, it is important to note that arguably most of what has been learned about the psychology of reading has been learned using the behavioral methods. However, it is also important to note that technology has rapidly expanded the arsenal of methods that reading researchers now have at their disposal. These new technologies allow researchers to make informed inferences about the mental processes that support reading by examining their neural correlates, the activity of the cortical systems that support cognition.

The oldest of these more recent methods entails the recording of the electrical currents generated by large ensembles of neurons that are engaged in the coordinated activity that occurs during reading. For example, as a participant names a sequence of words in the naming task, an *electroencephalogram* (*EEG*), or recording of the electrical activity of the participant's brain, is first recorded and then averaged across the individual responses to produce *event-related potentials*, or *ERPs* (for a review, see Handy, 2005). The ERPs that are collected from two different conditions (e.g., words read with vs. without normal parafoveal preview) can then be compared to make inferences about what happens during word identification (Antúñez et al., 2021). Because these recordings are analog and recorded continuously over time, ERPs have a fine temporal resolution, allowing researchers to examine how the brain's electrical activity changes over small intervals of time (e.g., milliseconds). The main drawback of this approach, however, is that the neural generators that give rise to the ERPs are difficult to localize. In other words, although the electrical activity is often recorded from a high number of electrodes (e.g., sixty-four) that are widely distributed across a participant's scalp, and although sophisticated mathematical techniques can be used to make inferences about the foci of the neural generators (Jatoi et al., 2014), the spatial resolution of these localization techniques is poor, often only allowing coarse inferences about the location of a neural generator (e.g., left vs. right cerebral hemisphere).

Fortunately, a few more recent brain-imaging methods have been developed to sidestep the problems of ERPs. The oldest of these methods, *positron emission tomography* (*PET*) and *functional magnetic resonance imaging* (*fMRI*), measure cortical activity indirectly, by measuring changes in blood flow that occur with increases in neural activity. With PET, these blood-flow changes are measured using radioactive tracers (for an introduction to this method, see Raichle, 1983). With fMRI, participants are placed in a strong magnetic field, so that the hydrogen atoms in the blood can be aligned with the magnetic field. With each off pulse of the magnetic field, the hydrogen atoms relax (i.e., spin randomly) and emit radio waves that are then detected by antennae situated around the participant's head (see Logothetis, 2003). Although both methods allow much better spatial resolution than ERPs, the resolution of fMRI is superior (typically a few cubic millimeters) and fMRI is less invasive because it does not require the injection of radioactive tracers. Where the two methods fare less well, however, is temporal resolution: PET measures brain activity within a given cortical region across several tens of seconds, whereas fMRI measures brain activity across several seconds.

The final method to be reviewed here, *magnetoencephalography* (MEG), is related to EEG in that it uses highly sensitive sensors to measure the magnetic induction that is generated by the post-synaptic potentials of neurons, with this measure of induction then being used to generate a composite image of the brain's electrical activity (see Baillet, 2017). This method is much less invasive than PET, with a spatial resolution comparable to fMRI but with a temporal resolution comparable to EEG. With all these strengths, one might ask why MEG is not used instead of the other brain-imaging methods? Apart from a few practical limitations (e.g., operating costs), one of the main limitations of MEG is that it is better suited to measuring the magnetic induction generated by neurons located near the surface of the brain rather than those generated by deep brain structures. A second limitation is that relative to both PET and fMRI, the nature of the signals being measured are complex and thus more poorly understood. For those reasons, although MEG is an extremely useful methodology, it is probably more useful to view MEG as complementary to (rather than superior to) the other brain-imaging methods.

With that brief introduction to the methods used in reading science, it is important to emphasize that *none* of the methods that have been mentioned in this section – behavioral or neurophysiological – are without limitations, and that *all* the methods have proven useful, especially when two or more methods converge to provide mutual support for some finding or conclusion. Our approach throughout the remainder of this book will therefore be to sample from these methods in a manner that allows us to cover key findings related to the psychology of reading, utilizing the insights afforded by each method.

1.3 Models of Reading

As indicated previously, a key marker of progress in understanding the psychology of reading is the fact that there have been several formally implemented theories or *models* of the core processes that occur in the minds of readers. Reichle (2021) provides a comprehensive review of many of these models, which for the purposes of facilitating their exposition, are grouped according to what the models attempt to simulate and explain:

1. the identification of printed words;
2. the syntactic and semantic processing that is required to combine words into representations of phrases and sentences;

3. the processing that is required to construct representations of more extended discourses (i.e., the meaning of two or more sentences);
4. how each of the aforementioned processes are coordinated with attention and the oculomotor system to determine when and where the eyes move during reading (i.e., eye-movement control in reading).

Because models of Chinese reading are discussed at length in upcoming chapters of this book, it is important to have a basic understanding of what computer models are, and why they are useful. For those reasons, two such models will be reviewed briefly here. Although these two models describe processes that are involved in the reading of English, the models are important for later discussions because the theoretical assumptions of the models have been borrowed in developing models of the reading of Chinese.

The first model to be reviewed here is the *interactive-activation* model that was first proposed by McClelland and Rumelhart (1981; Rumelhart & McClelland, 1982).⁴ This model, which is depicted in Figure 1.1, provides an example of an artificial neural network, and as such, consists of an interconnected network of *nodes*. These nodes are arranged in a hierarchy to represent the features or line segments of individual letters (in the bottom layer), letters (in the middle layer), and words (in the top layer). As shown, these nodes are interlinked by configurations of *excitatory* (arrows) and *inhibitory* (filled circles) connections that propagate activation among these nodes. Although this propagation of activity can be loosely conceptualized as corresponding to the propagation of neural activity among cortical areas that represent different types of lexical information (features, letters, and words), the interactive-activation model can also be viewed as an abstract description of the representations and algorithms that are engaged during word identification. Finally, it is the pattern of interconnections that is important for how the model functions. This pattern of interconnections is also what allows the model to explain the word-superiority effect (Reicher, 1969; Wheeler, 1970) discussed earlier.

As Figure 1.1 shows, the presence of a word causes the nodes corresponding to letter features (i.e., line segments corresponding to segments of the highly stylized font that is, for convenience, used to represent each letter) to become active. These letter features have specific locations, so that only

⁴ For a detailed description of the model including the equations that determine how excitatory and inhibitory activation is propagated among the different types of representational nodes, see either the original articles (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982) or Reichle (2021: 101–7).

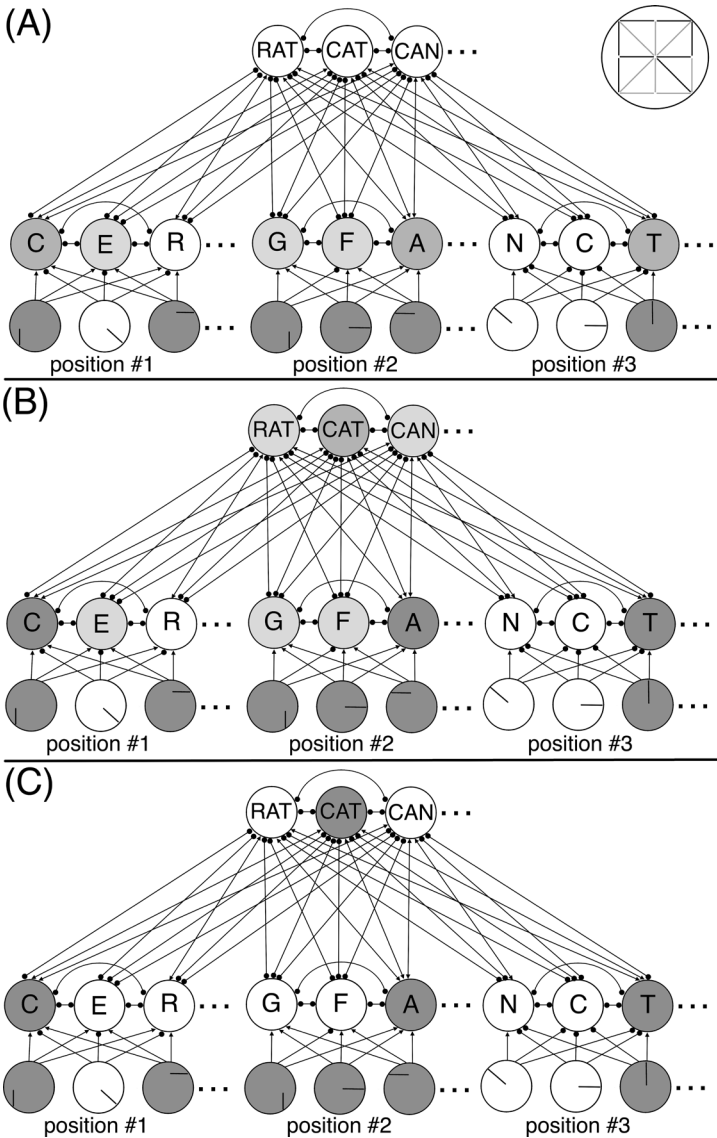


Figure 1.1 Schematic diagram of McClelland and Rumelhart's (1981) interactive-activation model of word identification

The nodes representing letter features, letters, and words are indicated, as are the excitatory (arrows) and inhibitory (circles) connections among nodes. Panels A–C show how the activation of nodes increase and decrease in their relative levels of activation (with darker gray representing more activation) in response to the word “cat” at three arbitrary points in time. The inset in the upper right of Panel A shows the full set of twenty features used in the feature nodes, with the dark gray indicating those features that would be active to represent the presence of a letter “R.”

Note: For the sake of clarity, this figure depicts only a small portion of the model.

one set of features, for example, can potentially become active in a word's first letter position, a second set in the word's second letter position, and so on. The letter feature activation then propagates to nodes representing individual letters. For example, upon being presented with the word "cat," the features corresponding to the horizontal and vertical line segments of the letter "c" in the first letter position will become active, which then send their activation to a node representing the letter "c" in the first letter position, but also to similar looking letters (e.g., "e") in the same position.

Across time, a set of letter nodes will become active, which then propagate their activation to words nodes containing those letters. As Figure 1.1 shows, the word "cat" will activate the nodes corresponding to its letters in the first through third letter positions, and as the activity of those nodes continues to ramp up, they will begin propagating their activation to word nodes containing at least some of those letters. As shown, the node for "cat" will become active, but so too (but to a lesser degree) will the nodes for words like "can" and "rat" because these words share some number of letters with "cat."

Finally, notice that, as the word nodes increase in their activation, the most active node will eventually come to suppress the others in a "winner take all" manner via the set of mutually inhibitory connections among the word nodes. This mutual inhibition is necessary to ensure that one and only one node will be identified at any given point in time. And while this is happening, notice that the words nodes also propagate their activation back to the letter nodes to which they are connected, allowing a well-activated word node to support the activation of its constituent letters in a mutually reinforcing manner. This "top down" propagation of activation is what allows the interactive-activation model to explain the word-superiority effect: Whereas a letter presented in isolation will only receive significant activation from its letter-feature nodes, a letter that is displayed in the context of a word will receive activation from both letter-feature nodes and the word node with which it is connected.

The interactive-activation model can explain several empirical findings besides the word-superiority effect. For example, because the word nodes have a "resting" or baseline level of activation that reflects how often the words that they represent have been encountered in printed text, common words require less time to activate than rare words, thereby providing an account of the *word-frequency effect*, or the finding that common words are typically identified more rapidly than less common words (Reingold et al., 2012; Schilling et al., 1998). It is also worth noting that these and other successes, along with the model's conceptual simplicity, have resulted in

it being highly influential in the development of other models of English (and as we shall see later, Chinese) reading. For example, the interactive-activation model is a core component of several other models of word identification (e.g., Coltheart et al., 2001; Davis, 2010; Grainger & Jacobs, 1996; Norris, 1994; Perry et al., 2007; Zorzi et al., 1998) and models of eye-movement control in reading (e.g., Reilly & Radach, 2003, 2006; Snell et al., 2018). Additionally, the more general notion of activation being propagated among a set of highly interactive nodes has been incorporated into models of sentence processing (e.g., Spivey & Tanenhaus, 1998) and discourse representation (e.g., Kintsch, 1998). Acknowledging this influence, let us now turn to a second example of a reading model.

This second example is the *E-Z Reader* model of eye-movement control during reading (Reichle et al., 1998, 2012; for a review, see Reichle, 2011). In contrast to the interactive-activation model, which provides a detailed or computationally explicit account of a single reading process, that of identifying printed words, *E-Z Reader* provides a high-level, more descriptive account of how several components of the mind work in a coordinated manner to determine when and where a reader's eye will move during reading. Figure 1.2 is a schematic diagram of the model.⁵ As shown, it consists of an early pre-attentive stage in which visual information is propagated from across the entire visual field to the mind, but with the fine-detailed features about words being used for lexical processing and the coarser features (e.g., about the locations and lengths of words) being used for saccadic programming. Each of these two processing "streams" will be described in turn.

As Figure 1.2 shows, lexical processing is completed in two successive stages. The first, *familiarity check* stage corresponds to a rapidly available sense of familiarity (e.g., like the recognition response in dual-process theories of memory; Yonelinas, 2002) that is used as a heuristic to "know" that lexical access is imminent, thus signaling the oculomotor system to start programming a saccade to move the eyes to the next word. The second stage of lexical processing, which corresponds to *lexical access*, then continues until the meaning and pronunciation of the word are available from memory. As shown, the completion of lexical access causes the focus of attention to shift to the next word, and the initiation of whatever post-lexical processing is required to integrate the meaning of the just-identified word into the representation of the sentence that is being

⁵ For a detailed description of the model, see Reichle (2011), Reichle et al. (2012), or Reichle (2021: 397–407).

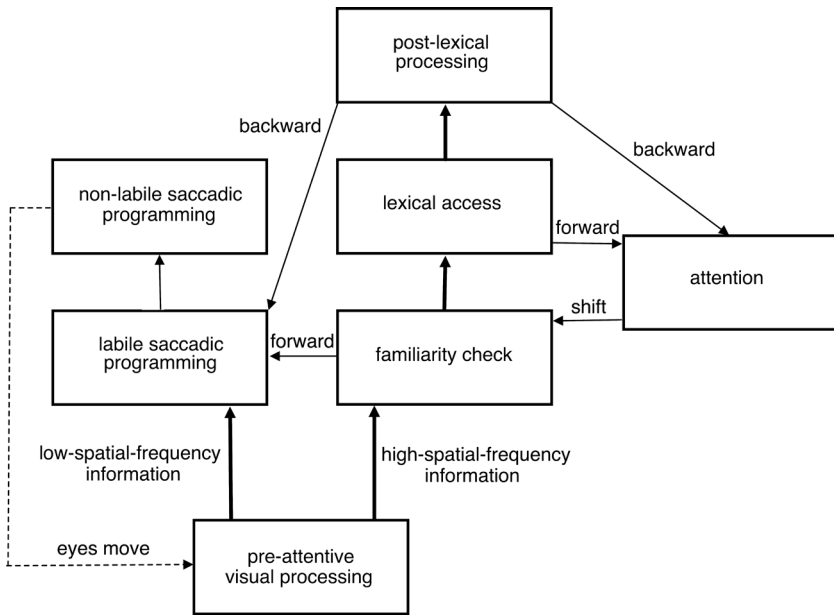


Figure 1.2 Schematic diagram of Reichle et al.'s (2012) E-Z Reader model of eye-movement control in reading

The boxes designate processes, the thick arrows indicate the propagation of information, and the thin arrows indicate the flow of control. The dashed arrow represents the actual movement of the eyes.

generated. As described so far, this part of the model instantiates the two core assumptions of E-Z Reader – that there is a dissociation between the events that trigger the movement of eyes (i.e., the familiarity check) versus attention (i.e., lexical access), and that attention is allocated in a strictly serial manner to support the processing and identification of only one word at any given time. Finally, according to the model, post-lexical processing occurs largely in the background on on-going lexical processing, only occasionally intervening if integration for some reason fails (e.g., the syntactic structure of a sentence is mis-parsed) or if integration is too slow (i.e., if word $N+1$ is identified before word N has been integrated). Either of these two situations can result in a pause or the triggering of an inter-word *regression* to move both the eyes and attention back to the source of integration difficulty.

The second processing “stream” in E-Z Reader is related to saccadic programming and execution. As Figure 1.2 shows, saccades are programmed

in two successive stages: an initial *labile stage* that can be canceled if another saccade is initiated, followed by a *non-labile stage* in which the saccade cannot be canceled. This distinction allows the model to explain why words are sometimes skipped (i.e., not fixated) during reading, as follows. Imagine a situation in which both the eyes and attention are on word N . In this situation, the completion of the familiarity check on word N will cause the oculomotor system to start programming a saccade to move the eyes to word $N+1$. Now imagine that, while this labile stage of programming is being completed, lexical access of word N completes, causing attention to shift to word $N+1$ and its lexical processing to begin. If the familiarity check of word $N+1$ then completes rapidly enough, it will trigger the initiation of a second labile program to move the eyes to word $N+2$, which then cancels the original program, causing word $N+1$ to be skipped. However, if the familiarity check of word $N+1$ completes more slowly, then the labile program to move the eyes to word $N+1$ will likely complete, initiating the non-labile stage of programming and thereby resulting in an obligatory fixation on word $N+1$.

Finally, although the saccades are always directed towards the centers of upcoming words (i.e., towards their *optimal-viewing position*; O'Regan, 1992) because this viewing location affords their efficient processing, there are two sources of saccadic error. The first is random and causes fixations to be normally distributed around their intended targets, but with the amount of deviation also increasing with the length of the intended saccade. The second type of error is systematic and causes saccades that are shorter/longer than some "preferred" length to over/undershoot their intended targets. Because both sources of error often result in fixations being in suboptimal viewing locations, the model also assumes that efference copies of the intended saccade can be used to quickly determine the size of the discrepancy, and to then rapidly initiate a corrective saccade to move the eyes closer to the originally intended target (i.e., the center of the word being processed). Together, these assumptions allow the model to explain why fixation landing-site distributions tend to be normal and centered near the middle of words (McConkie et al., 1988), and why fixations near either end of a word tend to be short in duration and more likely to be followed by a refixation on the word (Vitu et al., 2001).

More generally, the E-Z Reader model as described above has been used to simulate and understand many findings related to eye movements in reading (Reichle et al., 1998) and other reading-like experiments (Reichle et al., 2012; Veldre et al., 2023; for a review, see Reichle, 2011). And like the interactive-activation model (McClelland & Rumelhart, 1981) discussed

earlier, E-Z Reader has been influential, motivating a considerable amount of new empirical research (e.g., Inhoff et al., 2005; Pollatsek et al., 2006) and the development of several competitor models (e.g., Engbert et al., 2005; McDonald et al., 2005; Reilly & Radach, 2003, 2006; Schad & Engbert, 2012; Snell et al., 2018). More recently, the model has been “fleshed out” by embedding more computationally explicit models of word identification, sentence processing, and discourse representation within its framework to produce a computationally explicit account of reading in its entirety, *Über-Reader* (Reichle, 2021).

Finally, the two models that have been reviewed here, the interactive-activation model and E-Z Reader, are important for present purposes because they provide examples of the types of formal theories that have been developed to advance our understanding of the psychology of reading.⁶ This advancement occurs in two ways. First and foremost, the models provide useful summary descriptions of the main processes that are involved in reading, allowing researchers to think more concretely about what happens during reading, and to make predictions about what might happen in experimental situations. Such predictions are immensely useful for advancing the science of reading because they allow researchers to formulate precise tests that can be used to disconfirm one or more assumptions of a model, thereby allowing the model to be rejected in favor of other models, or for the faulty assumptions to be modified. (For discussion of how and why formal models are useful in psychology, see Hintzman, 1991.)

In the context of the remainder of this book, models like the two that have been described have a second important use. Because most reading models have been developed to explain the reading of languages that use alphabetic scripts, like English and German, the theoretical assumptions of those models may not be appropriate for understanding the reading of languages that use non-alphabetic writing systems, like Chinese. As we will argue later, these possible discrepancies are extremely interesting because they suggest one of two basic conclusions. The first is that the theoretical assumptions in question may simply be wrong, and that they must be replaced by assumptions that are general enough to explain the

⁶ Although both models have been formally implemented as computer programs, it is important to acknowledge that “formally implemented” is often a matter of degree, and that most models are implemented using some combination of mathematical equations, computer programs, and diagrams. That being said, less formally implemented models or verbal theories can also be important conceptual tools for both thinking and making predictions about the outcomes of experiments in new research domains, and for precisely that reason, a few examples of such theories are described in Chapters 3 and 4.

reading of, for example, English and Chinese. The second possible conclusion is that different assumptions may be required to explain the reading of English versus Chinese – that one set of assumptions may be necessary to understand the reading of one of the two languages, but either those assumptions are unnecessary or other assumptions are required to explain the reading of the other language.

Finally, given this brief discussion of why models are useful, one might ask about the process of adjudicating between two or more models. Or more generally, how are two or more models compared and evaluated? Although a complete answer to these questions can be extremely complicated (e.g., see Farrell & Lewandowsky, 2018), a short answer suffices for the purposes of this book. This short answer is that, with all else being equal, models that explain many empirical findings using a small number of theoretical assumptions are preferred to models that require many assumptions to explain just a few findings. Additional considerations that might be used in comparing and evaluating models might include: Do the models use assumptions that are consistent with what is known about either cognition or neuroscience more generally? And do the models generate predictions that are in some way novel or unexpected? After all, models are useful to the extent that they advance our understanding of some issue, and in relation to the psychology of reading, a useful model is one that provides a new insight into what might be happening in the mind of a reader as they convert the marks on a printed page into the rich and varied representations that are afforded by the capacity to read. Models of reading are useful because they can provide a window into how this capacity is possible.

1.4 Chapter Previews

This chapter has provided the basic information that might be required of someone without a strong background in cognitive psychology, linguistics, education, or one of their aligned disciplines to understand the remainder of this book. The next chapter will provide some additional background that may be especially useful for readers who lack an understanding of the Chinese languages and writing system, and the characteristics of the latter that are unique and that provide points of contrast for the research that has, to date, largely focused on the reading of alphabetic writing systems and European languages.

Chapters 3, 4, and 5 then comprise the core of the book, and as such, are organized similarly. For example, Chapter 3 will focus on lexical processing and word identification, beginning with a brief review of what has been

learned about these topics from the study of the reading of alphabetic writing systems (mostly English) using the experimental methods reviewed earlier in this chapter. The bulk of Chapter 3 will then focus on what has been learned about the processing and identification of characters and words in Chinese reading from experiments using the same methods. Chapter 3 will also review the models that have been developed to explain what is known about the identification of characters and words during Chinese reading.

Chapters 4 and 5 then continue using this same organizational approach, but with the former chapter focusing on skilled reading, and the latter focusing on the development of reading skill, its impairment (i.e., dyslexia), and what has been learned from cognitive neuroscience about the reading of Chinese. Because much of our own research has used eye tracking to study reading, much of the research on skilled reading that will be discussed in Chapter 4 is based on experiments that have also used this methodology. And although neuroscience methods of the type described earlier in this chapter have been used to study both the identification of isolated words and skilled reading, this research has been collectively relegated to Chapter 5 for the purpose of maintaining coherence. As each of these chapters will demonstrate, although there are consistencies in what has been learned about these topics across languages and writing systems as different as those used in the reading of English versus Chinese, there are also important differences – differences that are usually not afforded the recognition that they warrant, especially given the theoretical and practical implications that they likely have for our general understanding of reading.

Finally, Chapter 6 closes with a more explicit comparison of what has been learned about the reading of Chinese versus English (and other alphabetic writing systems), with particular emphasis on highlighting those points of contrast that might have important ramifications for the psychology of reading. This analysis will then be used to motivate a small set of outstanding questions – questions that, if answered, we believe might advance our basic understanding of what happens in the human mind when it is engaged in reading. These questions will then motivate our predictions about future research, and a few of the more basic challenges that remain to be addressed by future reading researchers. Our goal in doing all of this, however, is modest – if we are successful, we hope to provide a few “signposts” that might be useful to reading researchers who are interested in advancing the science of reading by studying what really is one of the most intriguing writing systems that was ever developed and that is still widely used today – that of written Chinese.