VIEWPOINT

INTERNATIONAL LITERACY ASSOCIATION

Following the Rules in an Unruly Writing System: The Cognitive Science of Learning to Read English

Devin M. Kearns, Matthew J. Cooper Borkenhagen

The authors describe the cognitive science of learning and how teachers can use it to improve reading instruction. They explain why some instructional strategies are essential and others detrimental to reading success.

Readers Do Not Learn Rules

A rule is a pattern that is so consistent that there are few exceptions. Many written languages were created so that letter patterns (we also use the term *spellings*) are matched with the same sounds (*phonemes*) in most words (e.g., Finnish, Korean). It can be said that they have "rules" in that we can describe succinctly how the letters relate to the sounds. English also appears to exhibit many rules: *M* is almost always pronounced /m/, for example. Explaining the system in terms of rules makes it easier for readers to pronounce printed words quickly and accurately, even if they have never seen these words before. For example, a Grade 1 student reading an origami guide would encounter the word *crease*. The student might have prior knowledge of the sounds made by the letters that comprise this word: *C* is pronounced /k/, *R* as /r/, and *ea* as /ē/. This knowledge might then be used to blend the sounds together to read *crease*. This apparently simple case is more complex when considering other possible sound-spelling patterns for these letters. Couldn't *C* say /s/ as in *cease*? Or /sh/ as in *ocean*? Couldn't *ea* say /ĕ/ as in *bread* or /ā/ as in *great*? Couldn't *S* say /z/ as in *tease*? And why doesn't the *e* at the end have a sound like it does in *recipe*? English may be an alphabetic language, but it seems like an unreliable one.

It does not just *seem* unreliable. It *is* unreliable—at least in part. Researchers describe English as a *quasiregular* writing system, meaning that the sound-spelling system has some very predictable "rules", but exhibits many exceptions. It follows from this that the reader cannot count on the rules to work 100% of the time—or even 90%, 80%, or 70% percent of the time. Rules are so unreliable that some researchers and educators prefer to call them *patterns*. We use the term *rules* because it helps us explain the predictable aspects of the writing system. However, describing the language as "rule governed" is problematic because (a) they do not apply in the vast majority of cases, and (b) at a mechanistic level reading development doesn't involve learning and applying rules that associate letters and sounds (more on this later).

Quasiregularity has at least one important consequence for education: It is hard to learn the English sound-spelling system. One source of evidence for this conclusion is that English-speaking children learn to read words more slowly than their peers in other European languages—all of which have greater regularity (Seymour et al., 2003).

The quasiregular nature of English is one reason that reading researchers and educators have long debated how much students need to learn about the sound-spelling system (see Seidenberg et al., 2020 for discussion). Some have suggested that the sound-spelling system is so confusing that it would be better to look at a couple letters and use context to guess at the word (and similar strategies; Goodman, 1967). This idea, while wrong, is appealing because it allows teachers to avoid having to explain the letters and sounds of the language. However, it does not lead to successful reading.

The purpose of this paper is to explain how readers learn the structure of the English sound-spelling system. The central idea presented here is that the learning mechanisms that underlie reading do not involve learning "rules" that associate print and speech even though rules help many learners understand the system by providing an onramp to remembering its most reliable patterns.

These facts have important implications for teaching students to read words early in the process of reading

Devin M. Kearns, is an associate professor, University of Connecticut, Storrs, CT, USA; email devin.kearns@uconn.edu. *Matthew J. Cooper Borkenhagen*, is an assistant professor, School of Teacher Education, Florida State University, The Florida Center for Reading Research, Tallahassee, FL, USA; email mcooperborkenhagen@fsu.edu.

 doi:10.1002/trtr.2299

acquisition. Teachers should indeed teach students to apply the most reliable and common sound-spelling rules to read words. Does this assertion contradict our point about the underlying nature of learning to read? No. Rules are a very useful tool for *decoding* (the process of linking letters to sounds to read unfamiliar words). Reading improves when students learn apparent rules and gradually understand the complexities that make reading *crease* a more manageable task. Teaching students the patterns in the sound-spelling system and applying these to read unfamiliar words is the essence of phonics instruction. Examples include teaching students that *A* say /ă/,[1] *A* says /m/, and other common patterns. Our goal here is to explain how the underlying learning mechanisms of reading do not involve rules, but that phonics is still helpful to gain knowledge about aspects of the English writing system. We also describe ways in which typical methods of phonics instruction are not themselves adequate to build a strong cognitive system for reading.

PAUSE AND PONDER

- English has a lot of sound-spelling *exceptions* exhibited in letter patterns that have multiple pronunciations. Does this mean that we should teach a very small number of regular letter patterns, or should we teach both the predictable patterns and their exceptions? What do you think?
- Some teachers like to teach students individual sound-spellings, like *T* says /t/ and *A* says /æ/. Other teachers focus on body-rime units like *AT* = /æt/. Which approach do you take?
- Reading is not hardwired in humans. When we read, we connect the letters, sounds, and meanings of words and use this knowledge to read sentences and texts. What happens when we try to identify a written word that we have never seen before—especially a word we can already say and define? Can we figure out the meaning from the letters? Do we need to connect the letters with their sounds? What are your assumptions about how this happens?

Learning, Cognition, and Reading

Learning Involves Gradual Differentiation Over Time

The observations about the learning process we've described thus far are true not only for reading, but of learning in general. To illustrate how learning works without rules, consider how learners learn about animals. At birth, a child is unable to differentiate animals at all. However, infants quickly acquire observations that allow them to understand important features that differentiate among animals they encounter in their environment. For example, some move and some do not, some have a soft furry exterior and others have scales, some have wings and can fly and some stay on the ground, and so on.

Over time as the child has experiences in the world, the child gradually builds knowledge about what makes each animal distinct from the others. The child begins to recognize the *features* the animals possess—various facts that are true about each (based on their visual properties, what they sound like, etc.). As the child learns about the features, their categories begin to emerge, as does the child's ability to reflect on (i.e., be *aware* of) these categories. The cat and cow will become more similar over time (they have legs, hair, and so on). Eventually, if someone asks a child what animals they know, they might report that cats, horses, and bears are animals, even if they have never said this out loud and even if no one ever explained this fact. The child will come to determine that a cat shares only some features with an ostrich, for example (it has legs but a different number; it has a tail but a different kind; see Figure 1 for a description of some cat features), but many features with a dog. A child who has accumulated a lot of knowledge might be able to tell you that a cat is more like a dog than a mouse because dogs are closer in size, also have whiskers, and can bare their pointy teeth.

Once the learner has used these features to develop a number of categories, the platypus inspires awe and consternation: It has hair like a cat, but a bill like a duck. To what category does this mystery creature belong? The child would likely be able to make observations but explain little, until you provided them with a useful descriptive tool: an explanation of the features that make a mammal. The platypus has fur, so it is a mammal.

The vast majority of this learning occurs without an explicit description of the features that make one animal distinct from another. Animal A has whiskers and paws with claws, preys on mice, does things like purr and hiss. Small children come to learn the associated label: "cat". A new animal will be called a cat if it shares a number of features with the other cats they've encountered. However, not all features of a new member of the cat category must match to assign the label: A Canadian Sphynx has no fur, but it has whiskers, pointy ears, and paws with claws, and does things like purr and hiss, so...it is a cat. In short, any concept we have labeled is not really a single thing. Cognitively speaking, the concept is characterized by a collection of features that *tend* to occur together. The idea of a *distributed representation* is a term sometimes used by researchers to describe the idea that any knowledge we have is distributed across features.

Figure 1
Common Features of a Cat

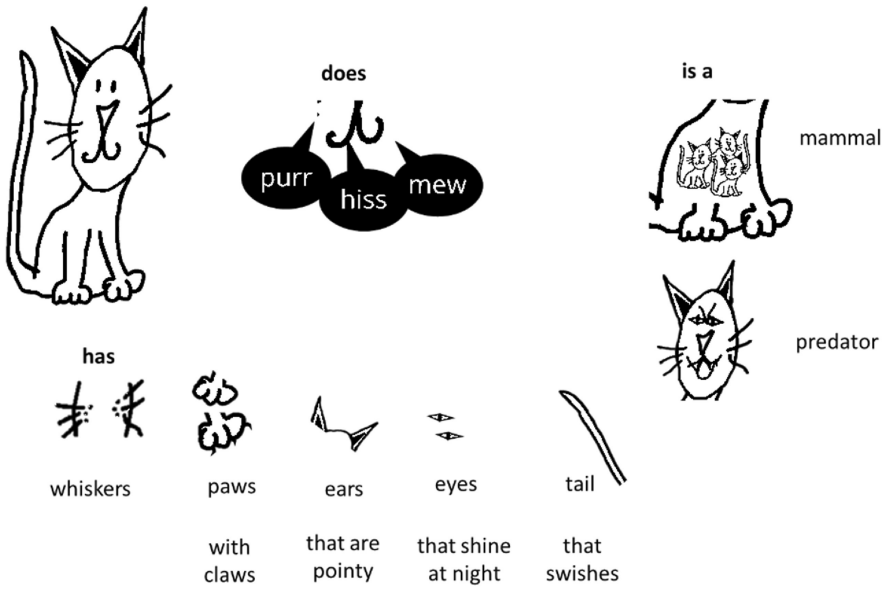

Information Comes in Through the Senses

Where does the child acquire this knowledge? From what information do they come to understand the relative differences among these elements of knowledge (e.g., different animals)? From their senses. They come to see, hear, touch, and even smell the differences between the different exemplars. Some of this knowledge remains implicit (something they know but cannot explain). Some knowledge is explicit (of which the child is consciously aware and can explain to you). Does the child need to know the rules of categorization to develop this knowledge? No. However, explanation can still help, especially if a teacher (at first a parent and later a school educator) wants to accelerate their learning. The teacher describes features that frequently differentiate types of animals to help the learner assign a new animal to the right category quickly.

By now, hopefully, the connection to early reading development is at least somewhat clear. A printed word is a collection of features. Cognitive scientists have shown that there are three types of features that are particularly important for reading and understanding printed words: what the word *looks* like (the letters, *orthography*), what the word *sounds* like (*phonology*), and what the word *means* (*semantics*). These three aspects of experience come from the senses.[2] The process of learning to connect the letters to the sounds and meanings is formalized in the *triangle model* of learning (Seidenberg & McClelland, 1989). This theoretical model of learning explains how reading development works. This model of learning also explains how learners acquire other knowledge about the world, including the labels and features of animals as described above. Put differently, the triangle model is designed to illustrate how learning works in general but is especially helpful for understanding reading. A depiction of the model is shown in Figure 2.[3]

A Model that Simulates the Learning Process

Researchers have tested the triangle model with computer simulations,[4] because it is difficult to inspect the human reading system (in the brain) directly. Evidence has shown that the model explains how learning works because the performance of the model is about the same as the performance of actual people at the end of the learning process. Researchers have shown that the pronunciation mistakes are similar to proficient adult readers, children who are learning to read, and individuals with reading difficulty at any age (e.g., Powell et al., 2006). This sounds complicated, but the process is straightforward. Below, we present a description to demystify the process.

The computer model is given an input representing the features of a word for its printed, spoken, or meaning form. For example, when reading *cat*, the input features are the letters. When the computer program presents a written word like *cat* to the model, a sequence of numbers representing the visual features of *C*, *A*, and *T* are provided as

Figure 2
A Diagram of the Triangle Model (Seidenberg & McClelland, 1989)

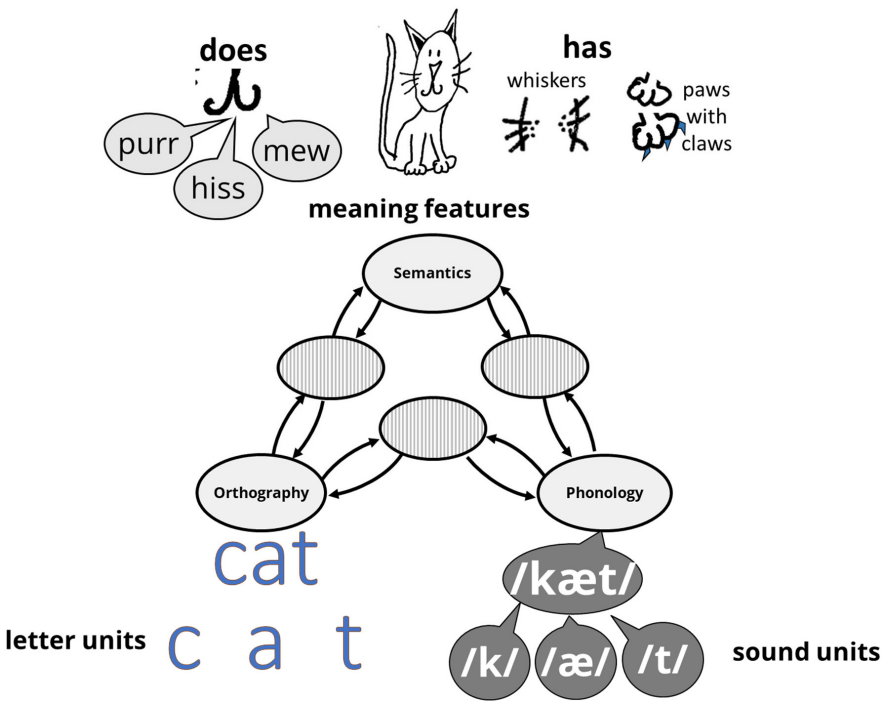

Note: Each corner of the triangle represents how information is processed during reading. The smaller ovals with vertical lines inside the corners and in between the arrows represent the internal representations of knowledge. These internal representations—called *hidden units*—are influenced by information flow through the arrows. These internal states represent the connections between the input (e.g., letters) and the output (e.g., sounds and/or meaning) that help differentiate one word from another. The connections between the inputs and outputs change with experience. If the reader makes a pronunciation error, a signal is sent back through the system to the letters, and the pattern of connections changes slightly. The model will perform better on that word on the next try. It will also perform better on other words that share some features. /æ/ represents the short *A* in the International Phonetic Alphabet (see Appendix A).

inputs. This numerical information flows through the model and a pronunciation (the output) is generated. The pronunciation is represented by the features of speech using the actions that occur in the mouth when pronouncing a sound. For the first sound, this means tapping the tongue against the back of the mouth (the /k/ sound), expelling a puff of air and then vibrating the vocal chords with the tongue and jaw in a particular position (the vowel), which is followed by a tap of the tongue behind the top teeth with a voiceless consonant (the /t/ sound). These actions are all represented by numbers that indicate the relevant features are on and the irrelevant ones are off for those three sounds. For the semantic features, there are features that represent the appearance and actions of a cat (as shown in Figure 1; purring, having whiskers, etc.). Letters, sounds, and meaning are collections of various features that define the words being learned. Features for a given word and for a given modality (i.e., orthography, phonology, or semantics) are spread across a number of different features that encode certain aspects of the word for that modality. As a result, we call these *distributed* feature representations.

Learning Happens by Adjusting the Internal Connections Between Inputs and Outputs

The description above explains how the learner takes in new information, but what about the information that the learner produces? When we see a word, the most common task is to produce its spoken form. This is our prediction of its pronunciation based on our knowledge at the time. When learning, the reader receives corrective feedback—like from a parent or teacher—if the reader gives the wrong pronunciation. This happens in the simulation too. The model predicts the pronunciation. The system compares the model's prediction to the actual answer (the true features for that word for that modality of output). Then, the network updates to increase the chances that next time that word will be pronounced in a way that is closer to the true pronunciation.

Knowledge is Shared Across Modalities (Senses)

What we learn through one modality (sense) can influence our knowledge associated with other modalities. In the case of reading, print and speech influence each other because they are related. In an alphabetic writing system like English (in which letters represent sounds), learners always improve their knowledge of letter–sound connections when they read a word correctly. In the presence of robust print and speech input (i.e., lots of co-occurring printed and spoken language, especially including common, useful words that show up in many diverse language contexts), we learn about the connections between those sensory experiences. The same is true for the other part of the triangle–meaning (though this aspect of experience may seem more abstract than the others). We very often use words in the presence of information about their meaning.[5] For example, we often name an object (e.g., a cat) in its presence. When readers encounter words they have never read before—but for which they already know the meaning—they instantly activate the meaning as soon as they determine the correct pronunciation.

All of this sensory information is taken in by the system, and our knowledge is updated based on feedback about what we produce and put into the world. The teacher (or some other source) might correct a spelling, pronunciation, or definition about a word based on the student's response. As a result, we eventually come to know the meaning associated with the letters as soon as we see it. This reflects the general idea about learning we have described: that we improve our knowledge across modalities whenever we determine the connections between information in one modality and another.

Categories are Internal Representations of Outside Information About the World

As learning to read happens gradually over time through experiences, categories form. Categories are mental representations of the features of a given input and output. In the model, the input categories are the letters, and the output features are the collection of meanings and the sounds in the word. This happens on many levels, large and small. The reader learns about the joint occurrence of letters and sounds based on the words they read and hear over time. For example, the letter *E* feature comes to be associated with the features representing the sounds /ĕ/ and /ē/ in different words (*bed* versus *me*). Likewise, when the *E* and *A* letter features are turned on next to each other, the features for /ē/ will likely be the output if following *B* and followed by *D* (as in *bead*). However, the /ĕ/ features will be produced in the context *H__D* (as in the word *head*). Even without teaching, the learner can acquire this knowledge and read words with *EA* words correctly. Put differently, the learner does not see a printed word only as a single letter in isolation. Rather, the learner visually processes other letters inside the word and uses all of the information from all of these letters to decide how to say the word[6].

Importantly, this process of learning is gradual. Knowledge emerges over time after enough experience, which progressively differentiates different elements of the environment from each other. For example, until the learner has lots of knowledge of the different contexts of *EA*, the pronunciation is not differentiated, and reading that pattern will contain errors. A reader would likely produce the most common sounds for the letter sequence, the long *E* sound /ē/ (like in *bead* and *each*). The reader will use that sound to read *meat*, *tread*, *great*, and *crease*. Experience progressively differentiates these words. The surrounding letters help the learner determine that the short-*E* pronunciation occurs more often in certain circumstances than others (see example with the letters *OLD* below). Think back to the animals: Until the child has enough experience, cats, wolves, and raccoons might be undifferentiated animals. When the features that distinguish animals are very slight, even adults struggle to differentiate them: Can you distinguish a rabbit from a hare? They share most of the same features and appear in basically all the same contexts.

Learning One Thing Affects Everything Else

As we learn, we accumulate vast knowledge about the features of words, namely those having to do with letters, sounds, and meaning. In addition, learners are not attending to just one feature of a printed word when we encounter it. We simultaneously perceive all of the visual information at once (for short words). For *cat*, this includes *C, A, T, CA, AT*, and *so on*. We also perceive positional information, like that the *C* is at the beginning of the word and the *T* at the end.

An important aspect of the learning process is that changes to the system affect other similar form—those that contain similar features to the word being processed. If the reader pronounces *cat* as *cot* (with a short o /ŏ/) and receives the feedback that the correct answer is *cat*, the reader changes how they will respond when presented with words containing the letter sequences in *cat* (the *A, AT, CA*, and *CAT*). It becomes more likely that the reader will pronounce *bad, rat, cap*, and *catnip* using the knowledge gained from cat. It might become less likely that the reader will pronounce *water* correctly—except that *A* in *WA* is frequently pronounced /ŏ/. If the reader had sufficient prior exposure to words with *WA*, the adjustment of *A* in

cat will have little impact on *WA* pronunciations. This is all how the tension between the frequency and consistency of patterns affect learning to read words.

The idea that words have distributed representations and that every learning experience changes the learner's entire network of knowledge distinguishes the triangle model from other models like orthographic mapping (e.g., Ehri, 2014). Orthographic mapping shares many similarities, particularly in that learning about sound-spelling units makes it easier for the reader to pronounce words correctly. An important difference is that orthographic mapping assumes that knowledge of the letters, sounds, and meaning of a given word are bound together separately for each word, but knowledge about structural properties are not shared across words that have similar structural features. The triangle model (and the simulations used to test whether this is true) involve no representations of individual words in this sense; knowledge exists in a network of features that are activated by an input (like seeing a word). The pronunciation and meaning associated with a visual word depends on the current state of the network, which has developed based on the reader's prior reading experience. Given that our reading knowledge involves our mental representation of features that occur together based on the words we've learned, rule-like structure emerges as we learn words that exhibit a particular pattern approximating the "rule". Importantly, this takes place regardless of whether the child has actually learned sound-spelling rules. A system of descriptive rules might help characterize this knowledge in some way, but does not reflect the underlying learning process. We discuss this further below. Before turning to that point, we will add that the orthographic mapping model overlaps enough that the instructional implications (described in the second part of this paper) may be similar.

Learners Appear to Learn Rules, But that is Not Really What is Happening

Given that learning to read involves learning about the features of words that occur together (letters, sounds, and meaning), "rules" should be thought of only as an abstraction of the most predictable patterns we come to learn. The reader comes to associate certain letters with certain sounds and certain meanings because they have been tightly connected in the reader's prior experience with words. Indeed, the reader can sometimes articulate a rule based on these experiences. This fact might seem far-fetched in the case of the most reliable patterns. Take for example the letter *B*, which almost always says /b/. Compare to patterns with less predictable pronunciations, though (like *EA* from earlier examples). The first author's last name can be used to illustrate this point. When strangers try to pronounce it for the first time, pronunciations fall into two categories: *Kearns* either rhymes with *learns*, or in such a way that the *EAR* is pronounced as in the word *ear*.[7] Which pronunciation the stranger uses depends entirely on their learning history, and the letters provide no clue to the standard pronunciation. The only way to know is to hear the author say it.[8] The point is that the reader's knowledge is shaped by experience. The role of experience is more obvious in the mystery of the *EAR* in *Kearns*, but it applies just as much to *B*—except that almost every encounter with *B* produces the same outcome and thus seems much more like a pattern exhibiting a rule.

To facilitate reading acquisition, teachers draw students' attention to the consistent patterns in the language. Teachers have intuited these patterns in the same way students have—over experience—except that teachers have a vast knowledge as expert readers that allows them to perceive many nuances not obvious to others. However, the quasiregular nature of English makes it challenging for the reader to depend on rules alone to identify many words. In addition, given the true nature of the underlying learning process, learning a phonic rule is unlikely to result in immediate, perfect application of this information. The reader will still require experience with printed words containing exemplars of any taught rule to develop knowledge in a way that will lead to lasting success with the associated language pattern.

This results in a clear tension that makes it difficult for students to understand the application of a rule: The rule expresses something that is only true in certain contexts and that might only make sense after a great deal of learning. As a result, rules can only be useful if the child has enough accumulated experience for the rule to be descriptive (e.g., they have to be able to distinguish between /b/ and /p/ as well as *B* and *P* and likely *Q* and *D*, etc.), and they have to have enough resources (cognitive and other) to be able to understand the description of the rule and its contexts (which will only be helpful for the most predictable patterns).

These Observations are True for Other Aspects of Language Knowledge

We have described the quasiregular nature of the English sound-spelling system and that our knowledge of language structure (including print) accumulates over time based on features and context. This is true for other aspects of sensory experience. There are even examples within other aspects of the writing system, including morphology, syllabic structure, and sentence structure. Morphological knowledge, for instance, also involves the gradual building up of knowledge about the relationship between spoken

language forms, meaning, and print. In this case, however, the segments of the language that emerge are larger, meaning bearing units, morphemes, rather than smaller letter-sound segments. That /d/, /əd/, and /t/ signal the past tense in the spoken words (/pād/ *paid*, /trŭstəd/ *trusted*, and /skrăcht/ *scratched*) is only true once sufficient knowledge of the sound, meaning, and print structure of the language has built up to support the differentiation of those forms. Likewise, our ability to determine the syllabic structure of the word *tepid*: /ˈtĕp əd/ (*TEP*-id) or /tĕ ˈpĭd/ (te-*PID*), depending on where you place the boundary, only emerges after vast experience with speaking words with longer phonological structure.[9] The spoken form of that word does not differentiate that boundary based on the auditory signal alone; it is only differentiated when we have experiences with many, many words with more than one syllable. So, as we move forward to talk about teaching, keep this in mind: much of what we are discussing here recapitulates itself in other aspects of language structure.

Why the Triangle Model's Theory of Learning and Cognition is Important for Teaching Reading

In this section, we describe how models of cognition are relevant for teaching. It is important to start with one essential fact that may be lost in our attempt to increase understanding of the link between learning cognition and early reading instruction: Phonics instruction is critically important. The reader should not mistake our recommendations to adjust some aspects of instruction as an argument against teaching sound-spellings. Phonics instruction has a long history of improving outcomes for readers, and it must be a major component of early reading instruction.

Given this context, here are some essential ideas based on the preceding section that have implications for how phonics instruction works:

- We learn gradually based on many experiences with the world—including with printed words.
- Given the way people learn, readers acquire an understanding of many features of the language, whether or not we teach them directly.
- Experiences with print help the sounds-for-reading processing system.

The way learning works has many implications for teaching, and we will describe some of the most important along with examples. Our goal is not to give the reader a specific strategy to use tomorrow. Our goal is to help reading teachers build lessons that are sensitive to these facts—whether that involves adjusting the way you teach a program or building new lessons on your own. As a practical matter, we always recommend adapting an existing program (especially one with known evidence of effectiveness) because of the incredible effort required to construct a comprehensive word recognition program—something we know from our own experience doing this. It is hard work for a team to accomplish this; it is extremely difficult for anyone to do on their own.

Idea 1 from the Triangle Model: We Learn Gradually Based on Many Experiences with the World—Including with Printed Words

Provide Students with Many Many Opportunities to Practice. This means that lessons should involve students and reading more words than the teacher. Student practice does not mean listening to another student read a word; practice means students reading aloud or silently themselves. If a student can respond at the same time as others in a choral response, that counts. Count how many times students practice reading a word (usually aloud to check the understanding of beginning readers but sometimes silently) every minute. How many times in 10 minutes?

Provide Opportunities to Practice that Allow Gradual Learning. Reading individual words "in isolation" is important for learning. Focusing on individual words is an important part of teaching word reading skills. What is even better is curating sets of words with similar and useful properties (Compton et al., 2014). However, extending practice to phrases, sentences, and texts is essential to support gradual learning as well.

Do Not Teach a Menu of Strategies that Use Conscious Processing—And Take Eyes Off Text. Teachers frequently provide students with a sort of menu of strategies—thinking that options allow students to choose something that works for them. This is not helpful because students do not learn to read words in fundamentally different ways from each other. In addition, while the menu may make sense to the teacher, strategies that require conscious processing (awareness how they are doing something) will reduce the reader's ability to learn the sound-spelling system in the way cognitive science suggest we do. Furthermore, the menu of strategies takes time away from practice. In word reading, the most problematic strategies are those that require lots of thinking and take students' eyes off text.

One prominent strategy is to suggest students look at the pictures to identify the word—sometimes called "Picture Power." The triangle model provides evidence and

a clear conclusion about this: Teaching students to look at the pictures to identify words is an fundamentally bad idea. It is perhaps the worst way to teach beginning readers to identify words.

English is an alphabetic language: The letters have the answers. English is quasiregular, so the answers are a little hard to figure out using rules. Nonetheless, the triangle model makes clear that the task of reading involves connecting information across modalities—but images associated with a printed word will not make it easier for a reader to pronounce it when encountering the word for the first time.[10] Spoken language knowledge is a critical component to the development of fast word recognition skills. Simply put, the beginning reader's job is to use the letters to read the words by connecting the letters first to speech. Telling the student to look anywhere else is not teaching students to read.

Almost as bad is a strategy that involves looking at the first few letters—usually the first consonants—and guessing the rest using sentence context. Readers might think that context is relevant; that is true. Just as pictures are an unreliable clue to the pronunciation of a word, so is the preceding context of the sentence. However, researchers have found that context serves as a confirmatory strategy that comes after making a decision about the pronunciation of the word using the letters (e.g., Gilbert et al., 2011; Share, 1995). If the reader says the wrong word and the teacher is not there to help, the reader might use context to correct the mistake. This decision must be based on analyzing all of the letters in order to produce a pronunciation. This is expressing the same idea as the point made before about linking sounds and meaning. An exactly correct pronunciation is essential for naming things. If the learner does not know the exact sounds in *cat*, it might be possible to confuse a cat with another word with similar features—like *bat*.

A third problematic strategy is to skip the word and return to it later. Students may occasionally experience enough difficulty reading a word that moving on is the only good option. However, skipping words often may cause readers to come to rely on skipping when they cannot immediately access pronounce the word. The alternative to the skip-it strategy is this: The teacher should (1) wait 5–10 seconds, (2) tell the student the word, (3) ask them to repeat it, and (4) allow them to continue reading. This exact wait time is not the important part; any sign of frustration or very slow deliberate decoding tells the teacher that this word is too difficult—even if the student might eventually decode it. Data strongly indicate that immediate corrective feedback supports learning (Epstein et al., 2002; Shute, 2008). Students will also probably benefit when the teacher draws little attention to the error—that is, the teacher moves on immediately after a short correction.

Some readers will know these three strategies as "Eagle Eye," "Lips the Fish" (who reminds the reader to "get your mouth ready"), and "Skippy Frog." The pictures are cute—there is even a Beanie Baby stuffed animal collection. However, these strategies are a serious detriment to reading success. Why? These strategies are entirely at odds with the way learning works.

Do Not Teach Students a Set of Strategies that Require Extensive High-Level Conscious Processing—Even if they Keep Student Eyes on Text. Another strategy-based approach to understanding the language is to use strategies for recognizing syllable types and teaching students syllable division. Before further explaining, it is important to note that some instructional programs that include these components have evidence of improving student reading achievement (e.g., Gersten et al., 2020). However, researchers have not shown that syllable practice is especially important to student reading success, and other instructional programs that also have evidence of improving student reading achievement do not use syllable strategies at all (e.g., Lovett et al., 2000). We mention this because syllable strategies themselves may not help with learning as much as their proponents would like.

There are two parts to these strategies. The first is to identify the type of syllable. By convention, there are six or seven types (opinions vary) that involve single-letter vowels (open and closed syllables) or another pattern: vowel-*R* (VR), vowel-consonant-*E* (VCE), vowel team (or digraph), and consonant-*LE* (CLE, or the "stable" final syllable). The seventh is a diphthong syllable used with special combinations (like for *OI* and *OU*), though not all agree it should be taught as a separate category. What is the reason not to identify these syllables? For VR, VCE, and vowel teams, there are frequent useful patterns in these words, but the data do not suggest they must be named as types of syllables. It might be just as helpful to point out that these are patterns found in words, not types of syllables.

The open and closed syllables are more challenging to characterize. It is true that single vowel letters at the end of a syllable often make the long sound (this is the open syllable) and that single vowel letters followed by a consonant make the short sound (this is a closed syllable). It is generally a good idea to teach about the type of sound based on its location within a syllable. It is probably not necessary to teach the terms *open* and *closed* to explain these patterns. It is adequate to call them *long vowel syllables* and *short vowel syllables*. It is the case that other patterns have long sound (like VCE), but those already have other names. It is much easier to rely on language students already know (the terms *short* and *long* are familiar

to them) than to teach something that has no real value outside of this context.

A syllable division strategy is very often linked to the syllable types. This strategy is designed to help students identify the vowel pronunciation of a single-letter vowel in a polysyllabic word. One rule is that the word *rabbit* can be divided using a VC|CV rule that makes the first syllable a closed syllable. The alternative, in the word *tiger*, would be the V|CV rule. Accordingly, the first syllable is open.

Readers can use these rules to pronounce many words: There are VC|CV words like *mammal* and *kitten*. The VC|CV pattern works reliably in two-syllable words. There are V|CV words like *emu* and *bison*. However, there are other V|CV words, like *lizard* and *camel*, that do not follow the pattern. The V|CV rule works more than half the time for two-syllable words (except for *E*, where the vowel is more likely to be short than long). The V|CV rule is highly unreliable in words with more than two syllables (Kearns, 2020).

Given this, are the rules worth teaching? The triangle model provides some supporting evidence. The core idea of syllable division associate involves associating letters and sounds. In this way, these rules are already much better than the eyes-off-text strategies. The triangle model provides less support in that the principle of gradual learning requires a massive amount of input, that is, lots of reading; the time taken to use the strategy reduces reading time. The gradual learning process is also mostly implicit (does not involve a conscious strategy). In this respect, the syllable strategies share a flaw with the eyes-off-text strategies, a need for concreted conscious effort. Whatever their flaws, the syllable strategies are still far better than the others described above because the goal is still to associate letters with sounds. At the same time, the triangle model shows why there is good reason to wonder about their value.

Idea 2 from the Triangle Model: Given the Way People Learn, Readers Learn About Multiple Aspects of the Language at Once, Whether Directly Taught or Learned Implicitly

Teachers Should Teach Phonics, Explicitly and Systematically. The fact that practice is important does not mean that phonics is not important. We have described the learning process in a way that does not involve rules. However, it is still important to explain sound-spelling associations that are reliable enough to support the ability to read unfamiliar words. Instruction on reliable sound-spelling rules is important because it helps shape the way the mind processes information. When we teach that *A* says /ă/, it helps solidify this idea in the reader's mind. As a result, upon seeing a word with a single *A*, we are more likely to say /ă/ than anything else. This is useful for beginning readers, especially when there are so many words with so many different patterns that figuring it out might take a very long time. By teaching the patterns, we can help readers learn the system faster than they otherwise would.

This is not a minor point. Despite English letters and sounds being linked, the system is still quasiregular. As a result, the important features might take a very long time to acquire. Many students do not learn the structure of the language fast enough to read texts with age-appropriate vocabulary and language demands. Phonics solves this problem: It helps the reader acquire the necessary information to apply to a greater range of unfamiliar words (e.g., Suggate, 2016).

Carefully designed and extensive phonics instruction is important in English—more so than in languages with consistent letter–sound patterns (e.g., Finnish). In those languages, learning is faster because the cross-modal (print-speech) information is more predictable, and the learner can build the sound-spelling system connections quickly and easily. English is harder because of the frustrating reality of quasiregularity. This fact requires that teachers draw conscious attention to many features of the reading system—including somewhat inconsistent ones.

To help students build the right connections, the teacher orders instruction to introduce more consistent rules first—and avoids exposure to inconsistent examples. The teacher will boldly declare, for example, that *C* says /k/, and admit successively that *C* also says /s/, and *K* also says /k/, and *CK* together say /k/. The teacher might clarify that the *CK* version of /k/ only comes at the end of a word, or in the middle when there is a suffix, and this spelling is used mostly when /k/ is preceded by a short vowel sound. But usually just in one-syllable words. There are even more unusual additions: *CH* can say /k/. But that is usually only with words of Greek origin. Notably, the teacher would not explain all of these facts about *C* and /k/; this just illustrates how many patterns could be considered in learning sound-spellings. However, the point remains that much of teaching English reading involves a string of confessions that there is more to the story than we said last time.

There is nothing wrong with gradually expanding students' understanding of the sound-spelling system. In doing this, the teacher is gradually shaping student knowledge, revealing more of its complexities to help them better use the letters to identify the sounds. It is much the same as helping a child differentiate cats from dogs by pointing out salient differences like their sounds (e.g., "what does a kitty say" or "look at the doggy's floppy ears"). We draw attention to more and more features that will help students succeed in linking letters to sounds. There is a limit to

the value of an explanation like this (a point we address below).

Students Will Acquire Understanding About Words that We Have Not Explained. No teacher (we hope) would teach students that the letters *GU* sound like /gw/ in words like *penguin*. Even without this information, students might read that word or *Uruguay* correctly the first time. The reason is that students are acquiring information about words even when we do not teach them. Teachers often find that students will begin to read words that do not follow patterns they have learned and have never been taught as whole words. The triangle model explains why. A reader looks at every part of the word at the same time (at least for shorter words), and the brain changes how it processes that word and what those letters could say just a little differently. Readers would probably never articulate why they can read these unfamiliar words with unexpected patterns, but they can still do it—because of these incremental changes based on their experiences.

Teachers Should Provide Instruction on Frequently Occurring Sound-Spellings. There are hundreds of possible sound-spelling patterns that we might describe by rule—from *GU* = /gw/ to *eaux* = /ō/ in *beaux* to *ae* = /ā/ in *aerospace*—and the teacher cannot—and should not—teach them all. The goal is to teach students enough sound-spelling rules to figure out new words on their own. The resulting logical question is this: How many sound-spellings should students learn? There is no easy answer. Learning more sound-spellings should result in the ability to read more words. If students were taught—and successfully learned—a large number of patterns, they would be able to pronounce many words correctly. However, the value of learning a pattern differs by the number of words it helps the reader pronounce correctly. In a database of 14,259 words present in texts for elementary-age children, only 81 occurred more than 100 times. This includes some that might surprise readers: The letter *A* making the "uh" sound (schwa, represented by /ə/ as in *about*) occurs in 1,583 words. The long *A* sound in *major* occurs in 629. A variety of other sound-spellings and their frequencies are given in Table 1.

In the table, the *OY* spelling occurs infrequently, despite that teachers frequently provide instruction about it. Should teachers explain the pattern and have students practice reading words containing it? Readers' own answers depend in part on their perspective about the benefit of teaching each sound-spelling relative to the number of words it helps students read. The long *E* pronunciation of *EIGH* (like in *Leigh*; see Table 1) would be a terrible sound-spelling if a teacher explained it—given that it applies only to two words. The long *I* pronunciation of *EIGH* (e.g., *height*) is similarly unhelpful. What about the long *A* pronunciation of *EIGH*? Is it worth teaching a sound-spelling for 27 words? Some will say yes.

Table 1
Sound-Spelling Frequencies

Sound-Spelling	Example	Number of words
ER = /er/	*her*	2,579
E = no sound	*house*	460
I = /ē/	*medium*	302
OW = /ō/	*town*	171
OW = /ou/	*now*	124
OO = /oo/	*boot*	175
OO = /uu/	*foot*	105
Y = /ī/	*fly*	75
OY = /oi/	*toy*	51
O = /u/	*prove*	41
EIGH = /ā/	*sleigh*	27
EIGH = /ī/	*height*	2
EIGH = /ē/	*Raleigh*	2
O = /ŭ/	*of*	1 (used 10,056 times per million words)

Some critics of extensive phonics instruction will point out that teaching *EIGH* and other infrequent and inconsistent patterns, like the six pronunciations of *OUGH* in English words, makes little sense. This is probably reasonable. However, these critics might also argue that much more frequent and consistent patterns should not be taught. Some have called for instruction on a limited set of only the very most frequent sound-spellings—the alphabet letters and a few more (e.g., Paris, 2005).

Limiting sound-spelling instruction to a very constrained set of patterns almost certainly underemphasizes the importance of teaching a variety of sound-spellings—but there are no rules for deciding the best number teach. If a sound-spelling occurs in more than 100 words, is that a good return on instructional investment? Probably (Vousden et al., 2011). We are inclined to say yes if the 100 words occurred frequently in texts themselves.[11] Whether 100 instances is a reasonable criterion for teaching a sound-spelling is not a decision with an obvious empirical answer. The answer probably depends on the context: How much information is adequate for a reader to continue building their knowledge of the quasiregular system without direct teaching? Some students will need more knowledge than others. The teacher will need to use

formative assessment data to make ongoing instructional adjustments.

Students Should Also Learn About Units Beyond the Sound-Spelling. Educators differ in their views about whether students should learn the pronunciations of larger sound-spelling units. Learning simulations implementing the triangle model show that learning does not require connecting one letter only to one sound in the way that phonics instruction would imply. Rather, learning involves processing all the information (features) of the word at once (for shorter words).

English sound-spelling patterns often differ depending on the surrounding letters. The table below shows the frequencies of the two pronunciations of *G*. The letter has the /g/ sound in 70% of words, but when *G* precedes *E*, *I*, or *Y*, it has the /ʤ/ sound 82% of the time. Although the *GE*, *GI*, and *GY* pattern is not perfect, it is relatively consistent and probably useful—applying to more than 200 words (Table 2).

Another example concerns the long *O* pronunciation of the letter *O*, as the table below shows. The pronunciation is not predictable for *O* alone. If the reader encounters *O*, it is just as likely to say the short sound as the schwa sound. If the reader attends to the *OL*, there is still not much consistency. However, when *O* is followed by *LD*, it almost always says /old/ and applies to 52 words (Table 3). The essential point is that the language contains some spoken language segments that involve attention to several letters. In some cases, it will benefit the learner for the teacher to draw attention to these larger patterns (Glazzard, 2017).

At the same time, larger patterns apply to fewer words. In the case of *OLD*, the pronunciation is consistent, but the frequency is relatively low. Is this pattern worth teaching? In some approaches to phonics instruction, educators focus primarily on phonogram learning. A phonogram includes the written vowel and trailing consonants within the same syllable. The idea is that phonograms are more reliable than individual sound-spellings and therefore more useful. *OLD* is a good example.

However, the trade-off concerns the number of words to which the phonogram applies. If we use the same 100-instance cut point for instructional worthiness, only 22 phonograms occur more than 100 times: *in, ist, ing, ic, us, is, um, ul, ip, an, on, and, it, am, ent, op, at, ect, ap, en, el, un*, many fewer than the 81 at the sound-spelling level. Focusing only or mostly on phonograms is unlikely to provide learners with enough information to make sense of the quasiregular system. Phonograms do appear to be useful to learners who have strong knowledge of the smaller sound-spelling units in words. Phonograms may be particularly useful for students after they build strong sound-spelling knowledge, but not before (Brown & Deavers, 1999; Schmalz et al., 2014). The general point is that teachers should consider the value of teaching sound-spelling patterns of different types and make thoughtful decisions about which and how many to teach based on student performance and what we've explained about learning here.

Table 2
Example of G Pronunciations Based on Multi-Letter Patterns

Complex pattern	Example	Number of words
G = /g/	go	891 (69%)
G = /j/	*gem*	383 (30%)
G = something else	*sign*	24 (1%)
G before *E, I, Y* = /g/	*get, gig*	56 (20%)
G before *E, I, Y* = /j/	*gem, gist, gym*	230 (80%)

Table 3
Consistency of the Pronunciation of O Based on Subsequent Letters

Pronunciation of O	Example	Frequency (% of all)
O = /ŏ/	*hot*	1,000 (37% of O)
O = /ō/	*go*	550 (19% of O)
O = /ə/	*together*	1,128 (34% of O)
O = something else	*to*	192 (7% of O)
OL = /ōl/	*gold*	105 (33% of OL)
OL = /ŏl/	*follow*	115 (36% of OL)
OL = something else	*police*	96 (30% of OL)
OLD = /ōld/	*old*	52 (95% of OLD)
OLD = /ŏld/	*Harold*	3 (5% of OLD)

Idea 3 from the Triangle Model: Experiences with Print Help the Sounds-for-Reading Processing System

This idea has become surprisingly controversial in recent years. Educators have questioned how much students need to link sounds to letters when they are first reading. Some programs include extensive phonological awareness practice involving manipulation of sounds without letters, and some have argued that excellent phonological awareness skills—especially with complex manipulations—must

precede word recognition instruction—or at least be the focus of much of student lessons.

The triangle model does not provide a clear answer concerning the amount of phonological awareness practice that will support reading skills. The part of the cognitive system that processes sound information works best when the reader has developed knowledge about sounds that frequently occur together (like /m/ before /p/, as in *camp*) and do not (like /m/ before /t/; *camt* is not a word and is difficult to pronounce). As a result, learners who understand well how sounds are connected may learn to read more quickly than those who do not. However, research does not clearly show that beginning readers and students with reading difficulty need extensive practice with sounds alone. The triangle model, along with extensive behavioral evidence, does show that the reading system improves as the learner has more experience connecting letters and sounds. As a result, building reading skills works best when reading involves letters. Teachers can structure activities to emphasize sound structure more. For example, they can use Elkonin boxes to make it clear that some letter combinations make one sound, and they can do activities that focus on manipulating sounds while using letter tiles (make *lip* into *lid*, make *lid* into *hid*, etc.). It is likely unnecessary to focus extensively on activities that only involve sounds. The triangle is very clear that reading is about linking letters to sounds. Once students begin to read, instruction should focus on using both.

Conclusion

The purpose of this article was to introduce readers to an important idea about the way people learn from their interactions with the environment through their senses. In reading, this process is described in the triangle model, simulations of which demonstrate how learning to read is a process that involves learning about the connections between the letters, sounds, and meaning of words. At a mechanistic level, learning is not simply a process of acquiring the "rules" of the writing system, though teaching students about reliable patterns is an important aspect of instruction. This model appears to overlap with the more widely known model of "orthographic mapping", despite differences in the assumptions about how the learner acquires knowledge of rule-like patterns. We have also described how teachers should provide instruction on rule-like patterns in the language because learning these patterns serves as an onramp to learning reliable properties of the language. A number of ways that this theory of learning can help teachers understand reading instruction have been discussed in a way that we hope will support instructional practices that align more closely with the learning processes described here. We hope this will inspire reading teachers to think critically about their phonics curriculum and adopt practices that align with these learning processes—whether it is a formal program or self-designed.

Conflict of Interest

We have no conflicts of interest to disclose.

ENDNOTES

[1] We use a sound code system often used in English dictionaries and books for teachers. However, this sound code is not used worldwide. The internationally accepted method for writing sounds is to use the International Phonetic Alphabet (IPA). We believe reading teachers should know about IPA. Our rationale is that this technical linguistic information allows us to interact with researchers and practitioners outside of our field—and they use these terms. The Appendix A contains a table for mapping the dictionary sound code to IPA. In this instance, /ă/ represents the short-*A* sound, /æ/ in IPA. Note that we are not recommending students learn IPA at any time in any phonics lesson.

[2] In psychological science, these are termed *perception* and *action*, corresponding approximately to information coming into the mind/body and information flowing out of the mind/ body. For our purposes here, *senses* will do. For more discussion on this point see Seidenberg et al. (2020).

[3] *Hidden unit* is a confusing term that we share only to help the reader encountering the term in the future. It is used in artificial intelligence systems to connect the inputs and outputs. It is essential to the idea that we do not have rules in that the features are actually represented by the strengths of the connections between the letter inputs, the hidden units, and the sound outputs.

[4] These simulations are commonly referred to as "computational models". The models are, essentially, computer programs that learn to map the printed words of the language to its spoken form and meaning using artificial neural networks.

[5] And note that a lot of this type of learning happens even before we start to learn about print. We come to reading development already having spoken language knowledge and knowledge about the world (though this is subject to variation across children).

[6] Though this takes on additional complexity for longer words and reading connected text, requiring eye movements to process successive chunks of visual information distributed across the page.

[7] *Kearns* rhymes with *learns*.

[8] The pronunciation of *Kearns* as containing *ear* is slightly more difficult to say, so this is the less likely pronunciation, but it is still quite common (strangers clearly unfamiliar with the name use the *ear* pronunciation 30–40% of the time).

[9] Note that this is precisely why phonological awareness (the self-awareness of this type of structure) emerges gradually over time. Our ability to isolate chunks of the language can only take place with enough accumulated knowledge of the relevant structure.

[10] A reader would be correct to point out that letters also represent meanings—the letters *C-A-T* represent a set of meaning features. The argument would be that the reader can directly connect letters and meaning by looking at pictures. The data from simulation research and the associated data from studies of people show that the connections with meaning only emerge after the reader has used letters to identify the pronunciation of the word.

[11] Here is an example where the number of words and the frequency of those words differs: O = /u/ as in into occurs in only 41 words, but it occurs in 41,973 words per million. Perhaps that changes the calculus. However, the calculus changes again if accounting for the

highest frequency words: to (30,785 instances per million words), do (4,091), into (2.761), etc. Without the 5 most frequent, O says /u/ only 924 times per million words.

REFERENCES

Brown, G. D., & Deavers, R. P. (1999). Units of analysis in nonword reading: Evidence from children and adults. *Journal of Experimental Child Psychology*, *73*(3), 208–242.

Compton, D. L., Miller, A. C., Elleman, A. M., & Steacy, L. M. (2014). Have we forsaken reading theory in the name of "quick fix" interventions for children with reading disability? *Scientific Studies of Reading*, *18*(1), 55–73.

Ehri, L. C. (2014). Orthographic mapping in the acquisition of sight word reading, spelling memory, and vocabulary learning. *Scientific Studies of Reading*, *18*(1), 5–21. https://doi.org/10.1080/10888438.2013.819356

Epstein, M. L., Lazarus, A. D., Calvano, T. B., Matthews, K. A., Hendel, R. A., Epstein, B. B., & Brosvic, G. M. (2002). Immediate feedback assessment technique promotes learning and corrects inaccurate first responses. *The Psychological Record*, *52*, 187–201.

Gersten, R., Haymond, K., Newman-Gonchar, R., Dimino, J., & Jayanthi, M. (2020). Meta-analysis of the impact of reading interventions for students in the primary grades. *Journal of Research on Educational Effectiveness*, *13*(2), 401–427.

Gilbert, J. K., Compton, D. L., & Kearns, D. M. (2011). Word and person effects on decoding accuracy: A new look at an old question. *Journal of Educational Psychology*, *103*(2), 489–507. https://doi.org/10.1037/a0023001

Glazzard, J. (2017). Assessing reading development through systematic synthetic phonics. *English in Education*, *51*(1), 44–57.

Goodman, K. (1967). Reading: A psycholinguistic guessing game. *Journal of the Reading Specialist*, *6*, 126–135.

Kearns, D. M. (2020). Does English have useful syllable division patterns? *Reading Research Quarterly*, *55*, S145–S160.

Lovett, M. W., Lacerenza, L., & Borden, S. L. (2000). Putting struggling readers on the PHAST track: A program to integrate phonological and strategy-based remedial reading instruction and maximize outcomes. *Journal of Learning Disabilities*, *33*(5), 458–476.

Paris, S. G. (2005). Reinterpreting the development of reading skills. *Reading Research Quarterly*, *40*(2), 184–202. https://doi.org/10.1598/RRQ.40.2.3

Powell, D., Plaut, D., & Funnell, E. (2006). Does the PMSP connectionist model of single word reading learn to read in the same way as a child? *Journal of Research in Reading*, *29*(2), 229–250. https://doi.org/10.1111/j.1467-9817.2006.00300.x

Schmalz, X., Marinus, E., Robidoux, S., Palethorpe, S., Castles, A., & Coltheart, M. (2014). Quantifying the reliance on different sublexical correspondences in German and English. *Journal of Cognitive Psychology*, *26*(8), 831–852.

Seidenberg, M. S., Cooper Borkenhagen, M., & Kearns, D. M. (2020). Lost in translation? Challenges in connecting reading science and educational practice. *Reading Research Quarterly*, *55*, S119–S130. https://doi.org/10.1002/rrq.341

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523–568. https://doi.org/10.1037//0033-295X.96.4.523

Seymour, P. H., Aro, M., Erskine, J. M., & Collaboration with COST Action A8 Network. (2003). Foundation literacy acquisition in European orthographies. *British Journal of Psychology*, *94*(2), 143–174. https://doi.org/10.1348/000712603321661859

Share, D. L. (1995). Phonological recoding and self-teaching: Sine qua non of reading acquisition. *Cognition*, *55*(2), 151–218. https://doi.org/10.1016/0010-0277(94)00645-2

Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research*, *78*(1), 153–189. https://doi.org/10.3102/0034654307313795

Suggate, S. P. (2016). A meta-analysis of the long-term effects of phonemic awareness, phonics, fluency, and reading comprehension interventions. *Journal of Learning Disabilities*, *49*(1), 77–96.

Vousden, J. I., Ellefson, M. R., Solity, J., & Chater, N. (2011). Simplifying reading: Applying the simplicity principle to reading. *Cognitive Science*, *35*(1), 34–78.

Appendix A

Table A1
IPA Representations for Sound Code

IPA symbol	Sound code symbol	Example	Type	Note
b	b	bat	consonant	
d	d	dog	consonant	
f	f	fish	consonant	
g	g	goat	consonant	
h	h	hat	consonant	
j	y	yes	consonant	
dʒ	j	badge	consonant	
k	k	cat	consonant	
l	l	lamp	consonant	
p	p	pen	consonant	
r	r	roll	consonant	IPA is sometimes written as ɹ, and some linguists would say this is the correct pronunciation in English.
ɹ	r	roll	consonant	
s	s	soup	consonant	
ʃ	sh	ship	consonant	
t	t	tap	consonant	
tʃ	ch	chat	consonant	
v	v	van	consonant	
w	w	wet	consonant	
j	y	few	consonant	
z	z	zoo	consonant	
ʒ	zh	vi<u>s</u>ion	consonant	
ð	<u>th</u>	that	consonant	
θ	th	thin	consonant	
aɪ	ī	bison	long vowel	This is technically a diphthong.
eɪ	ā	major	long vowel	IPA is sometimes written with *e* only. This is technically a diphthong.
iː	ē	see	long vowel	IPA is sometimes written as *i* only. The ː indicates lengthening; this is technically a diphthong.
oʊ	ō	boat	long vowel	IPA is sometimes written as *o* only.
ju	ū	unit	long vowel	This is technically a diphthong.
m	m	mouse	nasal consonant	
n	n	nose	nasal consonant	

(continued)

Table A1
IPA Representations for Sound Code (*continued*)

IPA symbol	Sound code symbol	Example	Type	Note
ŋ	ng	ring	nasal consonant	
ɔ	aw	hawk	other vowel	
ʊ	o͝o	foot	other vowel	There is no standard sound code for this vowel. It is sometimes written as *uu*.
u	oo	too	other vowel	When combined with *j*, it makes the long *U* sound.
aʊ	ou	bout	diphthong	A diphthong is a multi-phoneme pronunciation. The IPA is sometimes written as /æw/.
ɔɪ	oi	boil	diphthong	A diphthong is a multi-phoneme pronunciation. The sound code is sometimes written as /oy/.
ɑ	ŏ	pot	short vowel	
æ	ă	cat	short vowel	
ɛ	ĕ	bet	short vowel	
ɪ	ĭ	bit	short vowel	
ʌ	ŭ	but	short vowel	
ə	ə	about	unstressed vowel	This is commonly called schwa.
ɑɹ	ar	part	vowel-R	
ɔɹ	or	fort	vowel-R	
ɝ	er	her	vowel-R	The unstressed version of ɝ is ɚ, as in *mother*.
ks	ks	fox	multiple consonants	The sound of *X* is the combination of two consonants

Note: The term sound code describes a method of representing pronunciations that is often used in dictionaries and books for teachers. The International Phonetic Alphabet (IPA) is the standard way for presenting sounds that can be used across languages. We have used the sound code in the text because it is more familiar to readers, but we include this table to provide a guide for reading texts where the authors use IPA.