The Cognitive Foundations of Learning to Read: A Conceptual Framework for Teaching Beginning Reading

William E. Tunmer

Institute of Education

Massey University

The Cognitive Foundations of Learning to Read: A Conceptual Framework for Teaching Beginning Reading

Research indicates that the most effective beginning reading teachers share three characteristics in common. First, they have a broad understanding of the cognitive-developmental processes involved in learning to read, an understanding that guides the many instructional decisions that they make each day. Second, they can determine what beginning readers already know and what they still need to learn to become skilled readers. Third, they can provide students with targeted instruction that directly addresses their individual literacy learning needs.

The purpose of this learning module is to provide you with a graphically represented conceptual framework that will help you think critically about your instructional practice. We call this framework *The Cognitive Foundations of Learning to Read*. The module will first provide an overview of the major cognitive components upon which successful reading in an alphabetic writing system is based. It will then turn to the cognitive skills that underlie the ability to master these major components – in short, it will describe the cognitive structure of reading and then the foundation upon which that structure is built.

The module serves several uses for teachers. First, it offers a description of what the major skills in reading are, doing so at the highest level so as not to get lost in all the detail contained within these skills. Next, it provides an analysis of those skills that the would-be reader must master in order to find success. Note that these are not the skills that necessarily underlie skill in reading, rather they are those that underlie the development of skilled reading. Indeed, some of these skills, once learned as foundational for learning to read, have little to do with the skilled reading that is finally realized. The module will also help the practitioner understand that some skills must be acquired to fairly sophisticated levels before other skills can be acquired. This is critical, for trying to teach higher-order skills when the lower-order skills upon which they are based are weak, will be ineffective. Children taught this way will not be able to derive maximum benefit from instruction because of their insufficient level of development in the lower-order skills. The structure of the cognitive elements included in the framework also provides the basis for systematic assessment, the results of which can help you select appropriate instructional strategies for addressing the specific learning needs of beginning readers.

What is reading?

If the goal of literacy teaching is to help children learn to read, then perhaps a good place to start is to say what we mean by the human activity called reading. Reading can be defined as the process of extracting and constructing meaning from text. It is what you are currently doing as you derive meaning from the print on the page before you, print which to a pre-reader appears as a series of mysterious looking squiggles. An immediate difficulty with this definition, however, is that no reference is made to the language being read. Try reading this sentence: *Os resultados de dois estudos recentes opoiam essas proposicoes.* You most likely could not read it. Why not? It is because you are not familiar with the language being read. Being able to read English does not mean that you can read Portuguese!

Knowledge of the language being read is central to the reading process and without that knowledge reading could simply not take place. Reading is a language-based skill in which the reading process is grafted onto the listening process. From this perspective reading can be defined as the ability to translate from print to a form of code from which the reader can already derive meaning, namely, the reader's spoken language.

Although there are some differences between spoken and written language (e.g., speech exists temporarily, whereas print is relatively durable), comprehending text requires the full set of linguistic skills needed to comprehend spoken language. These include locating individual words in lexical memory (the mental dictionary where the meanings of the words that the reader or listener knows are stored) and determining the intended meaning of individual words (many of which have multiple meanings, e.g., "His *sentence* was awkwardly constructed" versus "His *sentence* was two years, three months"). Other linguistic skills required in comprehending speech and text include assigning appropriate syntactic structures to sentences, deriving meaning from individually structured sentences, and building meaningful discourse on the basis of the meanings assigned to individual sentences. These linguistic processes are described more fully in a later section of the module.

The Simple View of Reading (SVR)

Our proposed definition of reading suggests that the child's fundamental task in learning to read is to discover how print maps onto their existing spoken language. The process of learning to derive

meaning from print can therefore be adversely affected in one of two ways, or both: The child's spoken language system may be deficient in various ways, or the process by which print is connected to the child's spoken language system may be defective. These basic ideas are represented in the Simple View of Reading (SVR) model of the directly linked causes of individual differences in reading comprehension performance; the model provides an initial explanation for why some beginning readers perform well on reading comprehension measures while other children perform less well. The SVR model is fully incorporated into the Cognitive Foundations Framework described in the following sections.

The SVR model is widely supported in the scientific literature on learning to read, and has been very influential in the development of literacy policies in education systems throughout the world. In the UK, for example, the SVR model provided the conceptual framework underlying the wide range of recommendations included in the Rose report (*Independent Review of the Teaching of Early Reading: Final Report*), which was published by the Department of Education and Skills in 2006. The SVR model was adopted as the theoretical basis of the revised national curricular advice on the teaching of decoding and comprehension skills distributed to all schools in England.

The SVR model (graphically represented in Figure 1) proposes that at the highest level of analysis, the ability of beginning readers to derive meaning from text (*R*) is determined by their ability to recognize the words of text (*D*) and their ability to comprehend spoken language (*C*). The model further proposes that both *D* (word recognition) and *C* (oral language comprehension) are necessary and of equal importance in comprehending text (*R*). This fundamental idea can be represented in a simple equation, $R = D \times C$, where *D* and *C* range in values from 0 (no skill) to 1.0 (perfect skill). Thus, if word recognition ability is high but oral language comprehension skill is low, the child will be a poor reader (i.e., if $D = 1.0$ and $C = 0$, then $R = 1.0 \times 0 = 0$). If the opposite pattern occurs, where word recognition is low but oral language comprehension skill is high, the child will again be a poor reader (i.e., $D = 0$ and $C = 1.0$, then $R = 0 \times 1.0 = 0$). Stated simply, the process of extracting and constructing meaning from text (*R*) will be impaired if the child has trouble recognizing the words of text (*D*) and/or has trouble understanding the language being read (*C*).

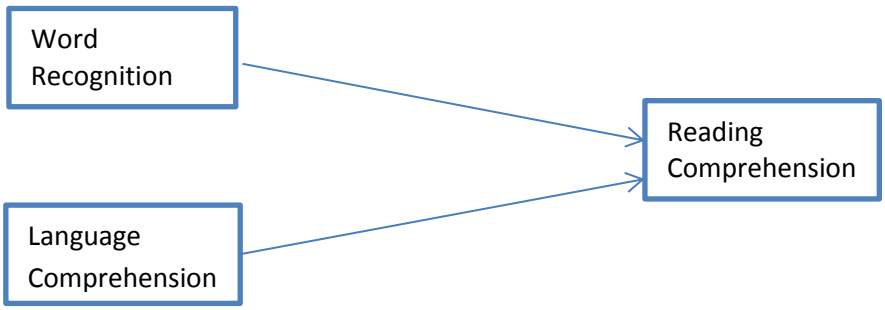

Figure 1. Components of the Simple View of Reading (SVR) model.

Preventing Negative Matthew Effects in Reading Development

As a model of the direct causes of reading performance differences, the SVR is not intended as a complete developmental theory of how children learn to read. Word recognition ability and oral language comprehension skill are themselves each dependent on the development of several other cognitive elements, all of which are represented in the Cognitive Foundations Framework that is discussed in the following sections. The important point to note here is that children who do not possess sufficient levels of mastery of these essential cognitive elements during the early stages of learning to read (and who are not provided with explicit instruction where needed to develop these competencies, especially those pertaining to the development of word recognition skills) will be forced to rely increasingly on ineffective literacy learning strategies, such as using picture cues, partial visual cues, and contextual guessing to identify unfamiliar words in text. The continued use of ineffective compensatory strategies inevitably leads to literacy learning difficulties. Children may rely on ineffective strategies to such an extent and for such a long period (years in some cases) that the strategies become entrenched and very difficult to "unlearn."

The consequences of continuing to rely on ineffective literacy learning strategies can be profound, as relatively small differences in essential reading-related cognitive skills during the early stages of formal reading instruction often develop into very large generalised differences in school-related skills and academic achievement. These downstream consequences are referred to as *Matthew effects*, or rich-get-richer and poor-get-poorer effects, after a passage from the gospel according to Matthew: "For unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken away even that which he hath."

Poor readers not only receive less practice in reading but soon begin to confront materials that are too difficult for them, which typically results in avoidance of reading, inattentive behaviour, low expectations of success, and withdrawal from literacy learning tasks (i.e., negative Matthew effects). As a consequence, they are prevented from taking advantage of the reciprocally facilitating relationships between growth in reading comprehension performance and the two constituent components of reading (i.e., *D* and *C*), a pattern referred to as positive Matthew effects (see Figure 2). As children become better readers, the amount and difficulty of the material they read increases. This in turn leads to greater practice opportunities for building fluency and facilitating implicit learning of letter-sound patterns (which improves *D*), and to growth in vocabulary knowledge, ability to comprehend more syntactically complex sentences, and development of richer and more elaborate knowledge bases (which improves *C*). Improvements in *D* and *C* promote further growth in *R* by enabling children to cope with more difficult materials.

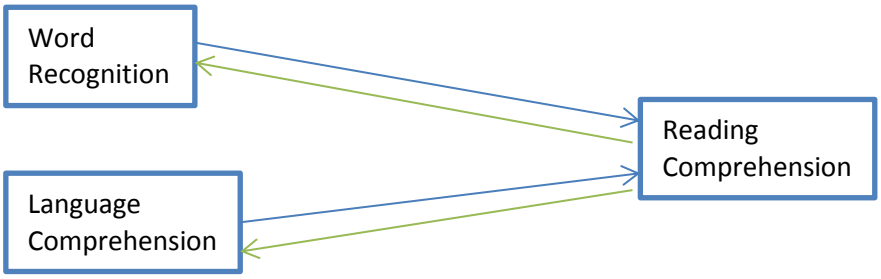

Figure 2. Reciprocally facilitating relationships (i.e., positive Matthew effects) between reading achievement (*R*) and both word recognition (*D*) and oral language comprehension (*C*).

As a result of repeated learning failures, many struggling readers also develop negative self-perceptions of ability and therefore do not try as hard as other students because of their low expectations of success and poor reading-related self-efficacy. For some of these children, especially boys, the sense of failure and feelings of frustration, coupled with the need to disguise their inability to perform literacy tasks, become so great that they begin to exhibit classroom behaviour problems. What began as relatively small differences in essential reading-related cognitive skills during the early stages of schooling soon develop into what has been described as a downward spiral of achievement deficits and negative motivational and behavioural spinoffs.

Given the nature of Matthew effects in reading, beginning literacy instruction that focuses on helping struggling readers to overcome weaknesses in the essential cognitive elements represented in the Cognitive Foundations Framework is a more effective teaching strategy than delaying action until substantial evidence of literacy learning difficulties has accumulated. The longer the delay in providing targeted, needs-based assistance, the greater the likelihood that reading problems will become severe and difficult to ameliorate. Any identified weaknesses in essential reading-related cognitive skills need to be addressed early and persistently by teachers and parents because they are necessary but not sufficient for the development of reading comprehension ability.

In the next two sections of this module we examine in detail the cognitive elements that underlie the development of the two constituent components of reading, oral language comprehension (*C*) and word recognition (*D*). These interrelated cognitive elements must be well developed in children for them to be successful at either comprehending language or recognizing printed words quickly and accurately.

Oral Language Comprehension (*C*)

In the SVR model, oral language comprehension is defined as the process by which, given lexical (i.e., word) information, sentences and discourses are interpreted. According to the SVR, children who have acquired oral language can understand age appropriate text when it is read aloud to them. If they become readers, then the comprehension these children experience when reading text is the same comprehension they would experience when listening to someone else read that text aloud to them.

What, then, are the processes involved in deriving meaning from speech? In spoken language comprehension words are built up from speech sounds, sentences are built up from words, and sets of interrelated propositions are built up from the propositions underlying individual sentences. The process of converting sound to meaning can be represented in terms of a model that specifies a set of interacting processors (i.e., mental mechanisms that operate in the brain) in which the output of each becomes the input to the next (see Figure 3).

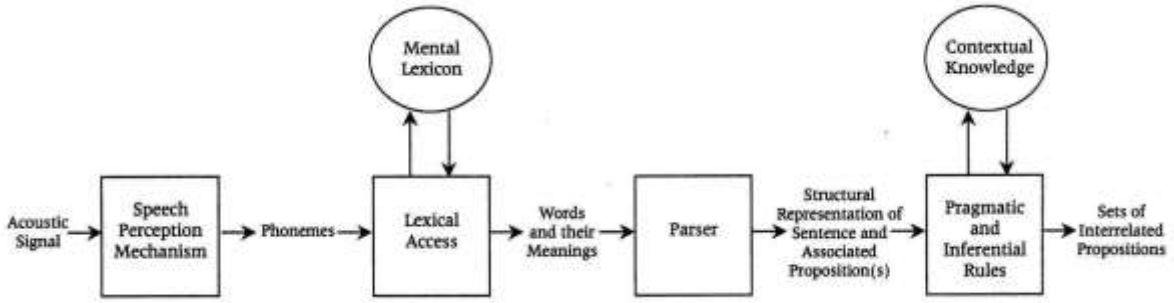

Figure 3. Model of processes and products involved in spoken language comprehension.

In oral language comprehension the acoustic signal (i.e., the sound waves produced by speech that arrive at the ear) is converted into a sequence of abstract units (phonemes) through a process referred to as speech perception. The process is highly complex, largely because there is no one-to-one correspondence between phonemes and segments of the acoustic signal. Information about phonemic content is transmitted in parallel; the information necessary for identifying a particular phoneme overlaps with that of another phoneme. You can get a sense of this by noting the difference in the position of your lips as you begin to say *tea* and *two*. But despite the parallel transmission of phonemic content, the speech perception mechanism produces as output a sequence of phonemes, or "phonemic tape" (see Figures 3). The phonemes then serve as the input to a lexical access mechanism that groups the phonemes and searches a mental lexicon (where the meanings of the words we know are stored in lexical memory) to find the meanings of the words in the utterance. The accurate identification of the individual phonemes in spoken words is therefore essential because the misclassification of a single phoneme can produce a word that is different in both meaning and syntactic category (e.g., *r*an vs *m*an). To comprehend spoken language, then, children must possess phonological knowledge; they must be able to distinguish the sounds (and the abstract units derived from those sounds) that make up the language.

The output of the lexical access mechanism shown in Figure 3 not only contains information about the objects, ideas, properties and actions to which words refer, but also information about how words can be combined with other words to form larger chunks of meaning. This information includes rules that govern the syntactic structures into which a word can enter (e.g., *The boy slept the bed* is not an acceptable sentence) and rules that place constraints on how words of different syntactic categories can be combined (e.g., *The cage slept* is not allowed). The meanings of some words (especially the function words of English, such as *the*, *a*, *on*) are expressed entirely in terms of rules specifying how the words can be combined with other words to form larger units of meaning. For example, try to define the meaning of the word *the*, which is the most frequently occurring word

in the English language. If you said that *the* is an article (or determiner), you only said what it is, not what it means. But you know that the phrase *the ball* means something different from *a ball*. Actually, the meaning of *the* (a quantifier) is very complex. Linguists have written entire papers about it!

Although the meaning of an utterance clearly depends on the meanings of the words it contains, it also depends on how the words are arranged, or parsed. To illustrate this point, go to the end of the next paragraph and begin reading the words in reverse order. You clearly know the meanings of the individual words but the paragraph makes little or no sense. Not only is the order of words important, but also how they are hierarchically grouped. Another processor (called the parser) must therefore take the words retrieved from the mental lexicon and build a structural representation of them, from which the utterance's (literal) meaning is derived (see Figure 3).

For example, the sentence *Flying hang gliders can be dangerous* has two entirely different meanings depending on how the words of the sentence are hierarchically organized (i.e., grouped). The first structural representation shown in Figure 4 represents the idea that the act of flying hang gliders (carried out by an unspecified person) can be dangerous to that person, whereas the second structural representation represents the idea that the physical objects themselves (i.e., flying hang gliders) can be dangerous to innocent bystanders (ignore the linguistic symbols that reflect the operation of rules that specify the structural representations of the two meanings of the sentence).

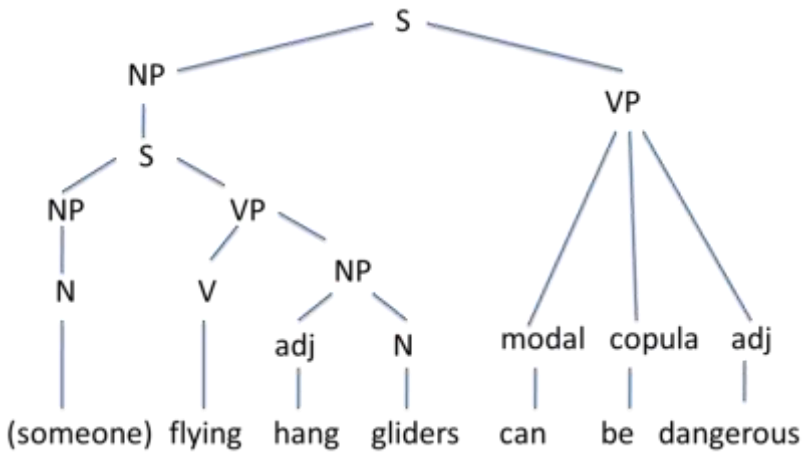

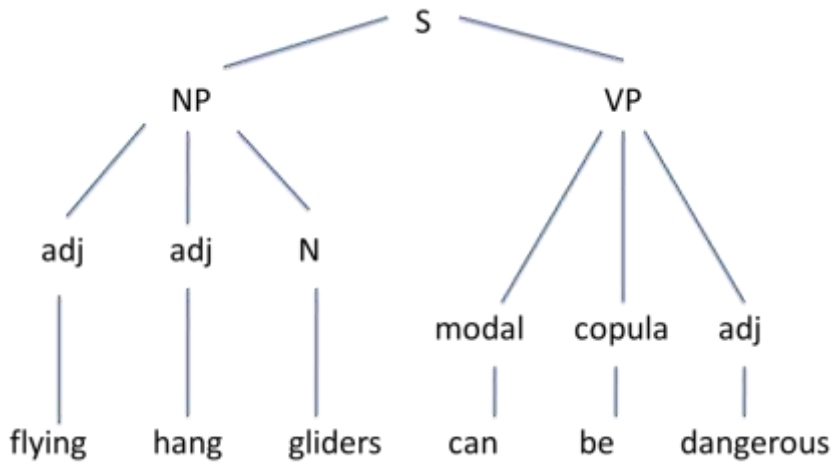

Figure 4. Structural representations for different meanings of the sentence *Flying hang gliders can be dangerous.*

The output of the lexical access mechanism that becomes the input to the parser (i.e., words and their meanings; see Figure 3), is placed in a temporary storage system called verbal working memory. Verbal working memory is often described as the "bottleneck" in the language processing system because the information that can be held in working memory is limited in both duration and capacity. To illustrate these two properties of verbal working memory, perform the following task: Quickly read the following list of words, immediately turn the page over when you have finished so that you can no longer see the words, and then attempt to recall all the words you just read in the order they appeared.

next
the
hair
the
to
with
girl
sat
boy
red
the

You may have retained some of the words by actively rehearsing them, but without rehearsal the words are quickly forgotten because they remain in working memory for only a brief period of time.

Also, due to the capacity limitations of working memory (usually fewer than 10 items), you probably found the task increasingly difficult as you progressed through the list. Now try performing the same task again with the following list of words (which are the same as before but re-arranged):

the
boy
sat
next
to
the
girl
with
the
red
hair

You no doubt found this task much easier to perform. This is because your parser grouped the words in the list into a smaller number of structural units which did not exceed the capacity limitations of working memory.

Because verbal working memory is limited in capacity and duration, the parser immediately begins attempting to build a structural representation of a new sentence as soon as words retrieved from lexical memory arrive as input. This is demonstrated by "garden path" sentences like the following:

> The large man weighed two thousand grapes.
>
> The shooting of the prince shocked his wife, since she thought he was a good marksman.

The element of surprise or confusion you may have experienced when you reached the end of these sentences occurred because the parser immediately began assigning a structural representation of each sentence that seemed probable given the content at the beginning of the sentence. However, the predicted structural representations turned out to be incorrect.

The processing limitations of verbal working memory have important implications for reading development. Word recognition processes that are inefficient and capacity draining make understanding text much more difficult for children. Readers with slow, non-automatic word recognition processes often forget the words they read at the beginning of a sentence by the time they reach the end. This makes it difficult, if not impossible, to determine the overall meaning of the sentence for two reasons: first, the earlier recognized words are no longer available, and second, due to the heavy expenditure of cognitive resources on word recognition, the remaining resources are inadequate for successful parsing and text integration processes.

As indicated in Figure 3 the propositions underlying individual sentences normally do not stand in isolation but are integrated into larger sets of interrelated propositions through the application of inferential and pragmatic rules that combine new information (i.e., the meaning of the sentence just processed) with previously existing information (i.e., prior knowledge of the situational context). For example, most listeners would not fully understand the sentence, *The cloth ripped but the haystack saved her*, unless they knew it was about a woman parachuting from an airplane, that is, unless they had comprehended the preceding discourse and made inferences that went beyond an interpretation based solely on linguistic knowledge. Similarly, inferences based on pragmatic rules for using language in social contexts enable listeners to understand the intended (as opposed to the literal) meaning of utterances. For example, if a parent said to their child, *The garbage is beginning to stink*, the child would infer from the situational context (i.e., the household distribution of responsibilities) that the parent means more than that a particular state of affairs exists in the world at a given point in time!

Background knowledge is the generic term used to refer to knowledge of the preceding discourse, prior knowledge activated by the developing meaning of the discourse, and knowledge of the situational context. To comprehend spoken language, children must have background knowledge that is relevant to what they are trying to understand. To understand written or spoken stories, they must have background knowledge that is related to the topic of the story. For example, children raised in New Zealand would more likely find it easier to understand a story about rugby than a story about baseball, whereas the opposite would be true of children raised in the United States. Understanding will also vary as a function of the kinds of cultural experiences children have had before entering school.

Inferences based on background knowledge can even influence the interpretation of single words in sentences that are structurally very similar. Consider the following sentences drawn from the research literature:

> The city council rejected the protesters' application for a parade permit because *they* feared violence.

> The city council rejected the protesters' application for a parade permit because *they* advocated violence.

In the first sentence the word *they* refers to the city council, whereas the referent for *they* in the second sentence is the protesters. In each case the referent for *they* is determined by inferences drawn from background knowledge.

Background knowledge also includes knowledge of schemas, which are knowledge structures that represent our knowledge of commonly occurring events and their relationships, such as going to a restaurant. However, simply having the relevant schema in long-term memory will not facilitate language comprehension unless it is activated. Consider the following example, which is also drawn from the research literature. Read the passage below and turn over the page when you have finished. Then attempt to write a detailed summary of what you just read.

> The procedure is actually quite simple. First you arrange items into different groups. Of course one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities that is the next step; otherwise, you are pretty well set. It is important not to overdo things. That is, it is better to do too few things at once than too many. In the short run this may not seem important but complications can easily arise. A mistake can be expensive as well.
>
> At first, the whole procedure will seem complicated. Soon, however, it will become just another facet of life. It is difficult to foresee any end to the necessity for this task in the immediate future, but then, one never can tell. After the procedure is completed one arranges the materials into different groups again. Then they can be put into their appropriate places. Eventually they will be used once more and the whole cycle will then have to be repeated. However, that is part of life.

Now imagine how much easier it would have been for you to carry out this task if you had been told at the outset that the passage was about washing clothes!

The role that background knowledge plays in language comprehension has important implications for teaching reading. When teachers ask children to listen to or read material, they can improve children's understanding of the material by providing or activating relevant background knowledge at the beginning of the lesson.

The relationship among the cognitive elements underpinning the development of the language comprehension component of the Cognitive Foundations framework is shown in Figure 5. This part of the framework draws from our description of the cognitive processes involved in deriving meaning from speech. Language comprehension is seen as a combination of literal comprehension based on linguistic knowledge (phonological, vocabulary/morphological, and syntactic) and inferential understanding based on background knowledge.

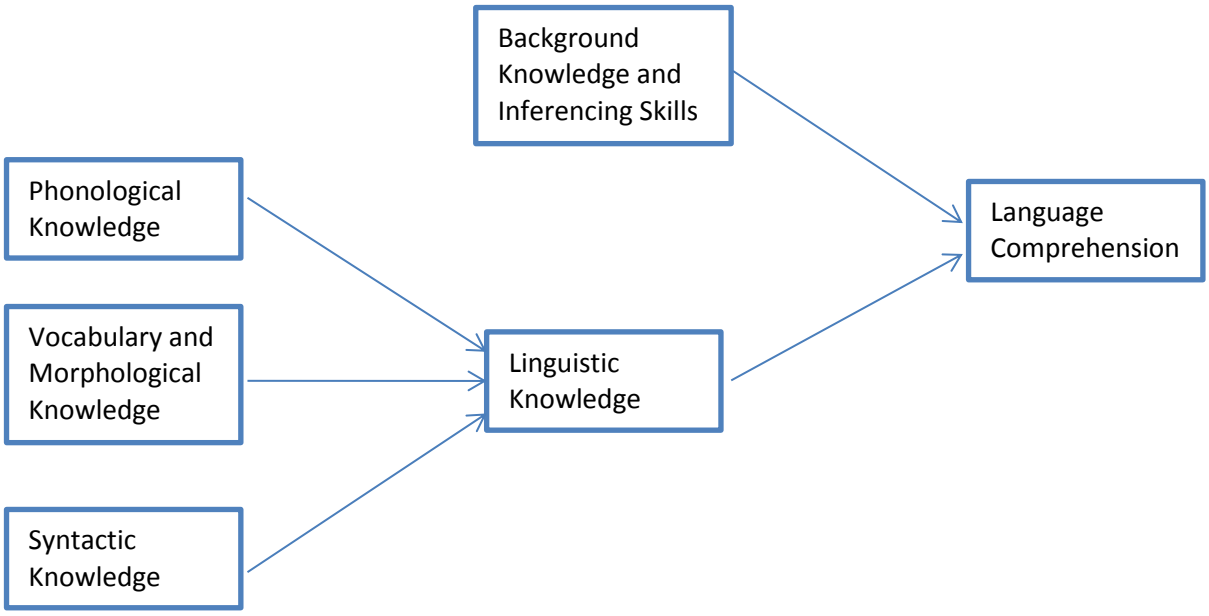

Figure 5: Relationship among the cognitive elements underpinning the development of the language comprehension component of the Cognitive Foundations framework.

Most children acquire their native language relatively easily through little more than exposure to a speech community during their first years of life. However, many children come from impoverished linguistic environments with limited pre-school exposure to the kinds of verbal interactions and language play activities in the home that promote the development of more advanced language skills, such as the ability to deal with the more formal, decontextualized (academic) language used in classrooms. For these children, explicit classroom instruction aimed at developing language skills is needed.

As noted previously, learning to read requires the full set of linguistic skills needed to comprehend spoken language. Weaknesses in the different levels of language functioning (as represented in Figure 3) would therefore be expected to result in different kinds of literacy learning difficulties. For example, children who have problems discriminating between different speech sounds because of a high-frequency hearing loss or deficits in auditory acuity due to otitis media (or "glue ear"), will encounter difficulty in analysing speech and relating it to print. Children with limited understanding of the meanings of the words of spoken language will be impaired in their ability to derive meaning from text. Children with poorly developed vocabulary knowledge will also have trouble identifying and assigning appropriate meanings to unknown printed words, especially partially decoded or irregularly spelled words, if the corresponding spoken words are not in their listening vocabulary. This in turn will limit the development of their alphabetic coding skills, as additional letter-sound relationships can be induced from the words that have been correctly identified.

Children with weaknesses in syntactic knowledge (i.e., implicit knowledge of rules specifying structural relationships in sentences) will have difficulty understanding written sentences and using the structural constraints of sentence context as an aid to identifying partially decoded words. Children who have problems in relating the meaning of each new sentence in spoken discourse to the meanings of the sentences that preceded it (i.e., discourse processing) will have difficulty comprehending and recalling written stories and passages. It is little wonder, then, that children who begin school with weaknesses in one or more of the subsystems of spoken language comprehension are much more likely to encounter problems in learning to read than children with age appropriate oral language skills.

Word Recognition (*D*)

Words are the essential elements of written language. Growth in reading comprehension depends on the ability to recognize the words of text accurately and quickly. Recognizing words accurately is important because the meaning of text depends crucially on the meanings of the words that make up text. Recognizing words quickly is important for the reasons mentioned in the preceding section; word recognition processes that are inefficient and capacity draining make understanding text very difficult as fewer cognitive resources are available for comprehension processes. Difficulties in recognizing words quickly occur when children attempt to identify most words by painstakingly sounding out and blending the letter sounds in words or laboriously using sentence context cues to guess the words. Acknowledging this, the SVR model defines word recognition as the ability to read words quickly and accurately; it is the ability to derive automatically a representation from printed input that allows access to the appropriate entry in the mental lexicon.

How does the development of automaticity in word recognition come about? For progress to occur in learning to read, beginning readers must acquire the ability to translate letters and letter patterns into phonological forms, called *alphabetic coding skill* (also referred to as *phonological decoding skill*). Alphabetic coding skill includes not only knowledge of correspondences between single letters or digraphs (e.g., *f*, *b*, *sh*, *oa*) and single phonemes (e.g., /f/, /b/, /sh/, /o/), correspondences between groups of letters (e.g., *tion*) and groups of phonemes (e.g., /*shun*/), and polyphonic spelling patterns (e.g., *ear* as in *bear* and *hear*; *own* as in *clown* and *flown*), but also knowledge of more complex conditional rules. These are rules whose application depends on position-specific constraints (e.g., the digraph *gh* at the beginning of words corresponds to /g/ as in *gherkin*, *ghetto*, *ghost*, and *ghastly*) or the presence of "marker" letters (e.g., the letter *e* indicates that the pronunciation of a vowel is long rather than short, as in *hop* versus *hope*; *tap* versus *tape*; *cut* versus *cute*; *bit* versus *bite*). The sounds of some letters are highly context-sensitive. For example, the letter *y* makes one sound in final position of 2-syllable words (e.g., *baby*, *happy*), another sound at the beginning of words (e.g., *yes*, *yell*, *yogurt*), and yet another sound in single open syllable words (e.g., *by*, *my*, *cry*).

Alphabetic coding skill also draws upon morphophonemic rules that speakers of English know implicitly through language acquisition, for example, that the morpheme for regular noun plural inflection (represented by the letter *s* in English orthography) is realized as /s/ when it follows an unvoiced consonant and as /z/ when it follows the voiced version of the same consonant, as shown

in the following examples (to understand the distinction between voiced and unvoiced consonants, place your fingers on your voice box and alternately say “ssss” and “zzzz”. You should feel your vocal chords vibrate when you say “zzzz” but not when you say “ssss”, even though place of articulation is the same for both sounds):

Unvoiced	Voiced
/s/	/z/
caps	cabs
cats	fads
sacks	sags
myths	lathes
muffs	dives

Similarly, the morpheme for verb past tense inflection (represented by the letters *ed* in English orthography) is realized as /t/ when it follows an unvoiced consonant and as /d/ when it follows the voiced version of the same consonant, as shown in the following examples:

Unvoiced	Voiced
/t/	/d/
reaped	grabbed
peeked	hugged
unearthed	seethed
puffed	loved
kissed	buzzed

For the bound morpheme meaning *not* (represented by the letters *in*), the pronunciation of the letter *n* reflects the phonological rule of spoken English specifying that within a word, a nasal consonant assumes the same place of articulation as the immediately following consonant, as shown in the following examples (the rule does not apply if the next phoneme is a vowel, as in *inoperable*):

indiscrete

implausible

inconceivable

In English orthography the spelling of the *in* morpheme in the word *implausible* is changed to *im* to reflect the change in pronunciation (which also occurs in the words *impossible* and *impractical*). Also, in line with the phonological rule governing the pronunciation of nasal consonants within words, the sound corresponding to *n* in the word *inconceivable* is /ng/, the sound at the end of words like *song* and *bang.*

The reason for drawing your attention to context-sensitive and morphemic spelling rules is to demonstrate that when examined at a deeper level, the English writing system is not nearly as irregular as often claimed. The different pronunciations of the letter *s* in the words *cats* and *dogs* reflects regularity, not irregularity, in English orthography.

Making use of letter-sound relationships to identify unfamiliar words is the basic mechanism for acquiring word-specific knowledge, including knowledge of irregularly spelled words. Taking advantage of the systematic mappings between subcomponents of written and spoken words enables beginning readers to identify unknown words, which in turn results in the formation of word-specific, sublexical connections between orthographic representations (i.e., the specific letter sequence of written words) and their corresponding phonological representations (i.e., the specific sound sequence of spoken words) in lexical memory, as shown in Figure 6. This process provides the basis for constructing the detailed orthographic representations required for the automatization of word recognition, thus freeing up cognitive resources for allocation to sentence comprehension and text integration processes. Correctly identifying words on the basis of letter-sound relationships a few times ultimately establishes their orthographic representations firmly in lexical memory from which additional letter-sound patterns can be induced without explicit instruction.

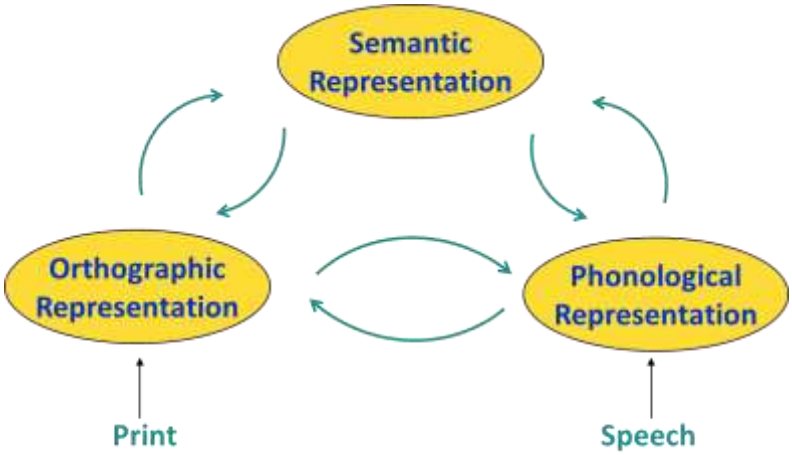

Figure 6. The development of word-specific, sublexical connections between the orthographic and phonological representations of words in lexical memory.

An alternative view of the development of word recognition skill to the one just described is that the ability to read evolves naturally and spontaneously out of children's pre-reading experiences with "environmental print", which refers to commonly occurring labels accompanied by context or logos, such as the word *stop* appearing on octagonally shaped signs at road intersections. According to this view, children are naturally predisposed to learn written language, so long as the emphasis is on the communication of meaning. Literacy teaching should therefore focus on meaning construction, not the abstract structural units that provide the basis for mapping print onto spoken language. Explicit instruction in word analysis activities is therefore downplayed or discouraged.

There are two major difficulties with this view of reading development. First, given that the world is awash in print, we would not expect that so few children learn to read before going to school, with those who do typically having received a considerable amount of encouragement and support in literacy-related activities in their home prior to school entry. Second, fast, accurate word recognition skills do not emerge from the kind of spontaneous word learning that results from exposure to environmental print, where children learn to "read the environment" rather than the word.

Children do, in fact, initially learn to read a few words through the natural strategy of associating some distinguishing feature of the printed word with its spoken counterpart as a whole. Any cue that will distinguish the word may be used by the child, such as something in the immediate environment in which the word appears (e.g., the golden arches of MacDonald's restaurants), a single letter or a matching pair of letters, the font in which the letters appear, the names of some of the letters if the child knows some letter names, or possibly a property of the whole word (e.g., its

colour, its length, or the resemblance of the whole word to a familiar object, such as the double humps in the word *camel*).

However, beginning readers who continue to learn to read words this way will face two major problems. First, although children will easily acquire a few words on the basis of visually distinct cues, the child's natural strategy of associating a familiar spoken word with some feature or attribute of the word's printed form will eventually break down. Each new word will become increasingly harder to acquire because of the difficulty of finding a unique cue to distinguish it from those that have already been learned. For example, if the child selects the presence of a squiggly character (i.e., the letter *s*) to recognize the word *stop*, then difficulties will arise when attempts are made to use this character to recognize the words *tops*, *post*, and *pots*. Notice that in this example the selection of *any* character as a distinguishing cue, such as the cross-shaped character (i.e., the letter *t*), will fail. Beginning readers will make an ever-increasing number of errors and become confused and frustrated unless they discover or are led to discover an alternative strategy for establishing the relationship between the written and spoken forms of language.

Second, the strategy of selective association based on distinctive visual cues is developmentally limiting because it is not generative; it does not provide a means for identifying words not seen before. This is an important consideration because most of the words that beginning readers encounter in print are novel. Beginning reading materials typically employ upwards of 1,500 words, each of which must be encountered a first time. Moreover, when new words appear in print they do not suddenly begin appearing with great frequency (e.g., as in a story about the *eruption* of a *volcano*). Beginning readers are continually encountering words that they have not seen before and may not set eyes on again for some time.

Progress in learning to read can only occur if children abandon the more natural non-analytic strategy of using partial visual cues to recognize words. They must instead learn to develop analytic links between print and speech by means of fully analytic processing that requires an explicit and conscious awareness that letters and letter combinations are used to represent the phonemes of spoken words, referred to as the *alphabetic principle*. For example, the child who writes the word *colour* as KLR clearly grasps the alphabetic principle.

For beginning readers who continue to rely mostly on partial visual cues supported by contextual guessing at the expense of phonological information, there is little interaction between the

subcomponents of written and spoken words. To understand what learning to read this way would be like, imagine a writing system in which each spoken word of English is represented by an arbitrary string of numbers ("cat" = 47937, "and" = 35690). Without the support of spelling-to-sound connections, the word recognition system would quickly become overwhelmed. The process of learning to read would be roughly equivalent to learning to recognize quickly and accurately 20,000 telephone numbers! A similar situation occurs when learning to read in a nonalphabetic orthography like Japanese Kanji, which is based on borrowed or modified Chinese logographs. It takes 10 to 12 years of devoted study to learn to recognize 1,000 to 2,000 logographs, whereas the average high school student who has learned to read an alphabetic orthography can recognize quickly and accurately 25,000 words or more. The word recognition skills of children who do not make use of letter-sound relationships in learning to read English will remain relatively weak because they do not develop as rich a network of sublexical connections between orthographic and phonological representations in lexical memory as do normally developing readers. Because of their inefficient and capacity draining word recognition skills, these children will experience progressive deterioration in their rate of reading comprehension development as they grow older.

Figure 7 represents the transition from analytic to automatic word recognition. Initially children acquire and make use of analytic links between print and speech (i.e., alphabetic coding skill) to identify unknown words. After words are successfully identified a few times, automaticity in word recognition is eventually achieved based on the development of word-specific, sublexical connections between orthographic and phonological representations of words in lexical memory (see Figure 6). Support for this conceptualization of reading development comes from recent neurological research using functional magnetic resonance imaging (fMRI) to study changes in the anatomical distribution of neurophysiological activity in the brain as children progress from unskilled to skilled readers. The general picture that has emerged from this research is that typically developing readers initially rely primarily on a system in the parieto-temporal area of the left cerebral hemisphere to identify words in a slow, analytic manner by taking them apart and linking their letters to their sounds. However, as the children's reading skills develop, a system in the occipital-temporal area of the left hemisphere gradually takes over. This system is described as the express pathway to reading, the "word form" area of the brain where printed words come to be recognized very rapidly on sight.

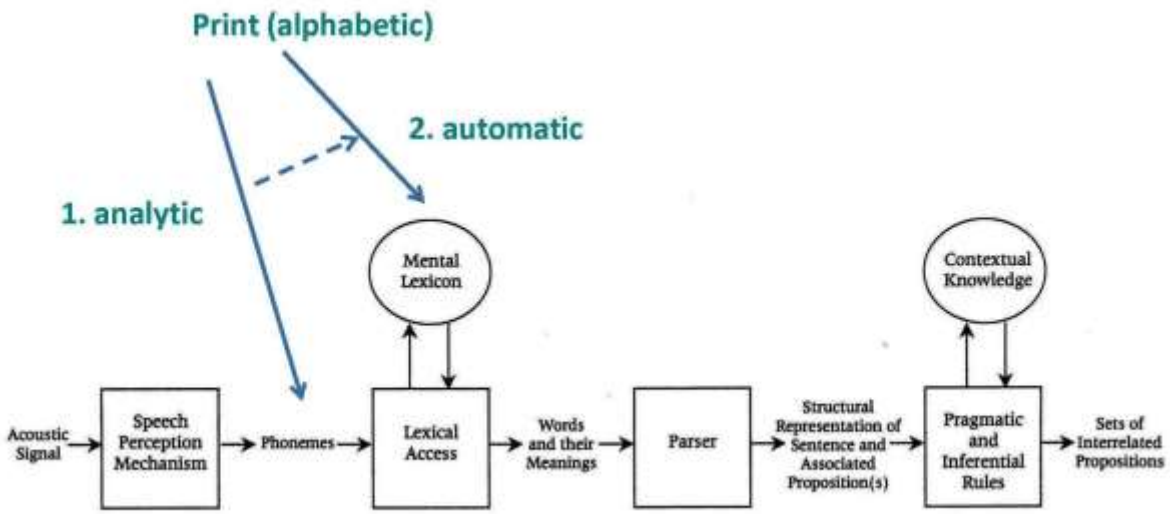

Figure 7. The transition from analytic to automatic processing of words in text.

In addition to knowledge of the alphabetic principle, *letter knowledge* and *phonemic awareness* are two other cognitive elements that need to be well developed before beginning readers can acquire alphabetic coding skill. *Letter knowledge* is the ability to recognize and manipulate the letters of the alphabet, including letters in different fonts and type case. Children should be able to discriminate each letter of the alphabet from all the others, which is usually accomplished by learning letter names. Familiarity with the letters of the alphabet is essential for developing alphabetic coding skill, where individual letters and letter patterns are mapped onto the sounds of spoken English.

Letter knowledge also contributes to beginning reading achievement in three other ways. First, letter-name knowledge serves as a bridge towards understanding the alphabetic principle, as reflected in children's invented spellings (e.g., *da* for *day*, *bl* for *bell*), where the names of letters are used to represent sounds in words. Second, letter knowledge acts as a precursor to alphabetic coding skill because the names of most letters contain the phoneme to which the letter normally refers. For example, the first phoneme of the name of the letter *b* (the phonological representation for which is /bi/) is /b/. Third, letter-name knowledge facilitates the development of phonemic awareness, especially when children are exposed to alphabet books and games that increase knowledge of letter names and their relation to sounds in words (e.g., "s" for snake). Most likely for these reasons letter knowledge is one of the best predictors of beginning reading achievement.

To discover relationships between spelling patterns and sound patterns, children must develop *phonemic awareness*, the ability to reflect on and manipulate the phonemic elements of spoken language. The phonemically aware child knows that the spoken word "feet" has three sounds, that

"eat" is what remains when the first sound (e.g., /f/) of "feet" is deleted, and that "fee" is what remains when the final sound (e.g., /t/) is deleted. Children who experience ongoing difficulties in analysing spoken words into their constituent phonemic elements will not be able to fully grasp the alphabetic principle and discover spelling-to-sound relationships.

Many beginning readers find it extraordinarily difficult to detect phonemic sequences in spoken words, even though they are clearly capable of discriminating between speech sounds and using phonemic contrasts to signal meaning differences. Using a phonemic contrast to signal a meaning difference (such as saying "pig" rather than "big" when referring to the farm animal), which is done intuitively and at a subconscious level, is not the same as consciously reflecting on and manipulating the phonemic elements of speech. Gaining conscious access to phonemic segments is much more difficult for children because, as noted previously, there is no simple physical basis for recognizing phonemes in speech. Children must develop the (metalinguistic) ability to perform cognitive operations on the products of the mental mechanism responsible for converting the speech signal into a sequence of phonemes (see Figures 3 and 7).

These considerations provide a likely explanation for why many children who have begun formal reading instruction fail to benefit from either letter-name knowledge or letter-sound knowledge in learning to recognize words. Because there is no one-to-one correspondence between phonemes and segments of the acoustic signal, it is not possible to pronounce in isolation the sound corresponding to most phonemes. Consequently, the beginning reading strategy of simply "sounding out" a printed word like *bag* will result in "buh ah guh", a nonsense word comprising three syllables. Most letter sounds and letter names are only imprecise physical analogues of the phonemes in spoken words. Whether children learn to associate the sound "buh" or the name "bee" or both with the letter *b*, they must still be able to segment the sound or name to make the connection between *b* and the abstract phoneme /b/, which cannot be pronounced in isolation (see Figure 8). Teaching activities designed to promote the development of phonemic awareness in beginning readers have been found to be very helpful, especially when combined with letter-sound training.

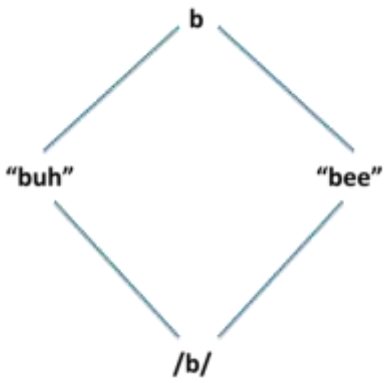

Figure 8. Phonemic awareness is required to segment the letter sound ("buh") or letter name ("bee") to make the connection between the letter *b* and the phoneme /b/.

The relationship among the cognitive elements that underpin the development of the word recognition component of the Cognitive Foundations framework is shown in Figure 9.

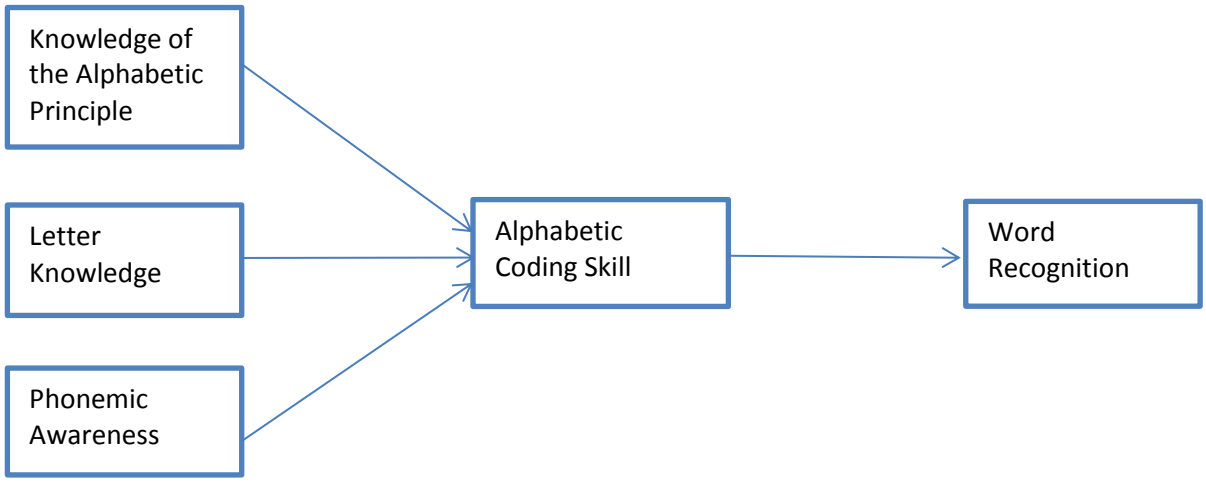

Figure 9. Relationship among the cognitive elements underpinning the development of the word recognition component of the Cognitive Foundations framework.

Facilitating the Development of Alphabetic Coding Skill

Assuming that the cognitive elements of phonemic awareness, letter knowledge, and knowledge of the alphabetic principle are reasonably well established, what are the best teaching strategies for

facilitating the development of alphabetic coding skill in beginning readers? Traditional phonics programs suffer from two major shortcomings. First, they tend to be strongly teacher-centred and have curricula that are typically rigid, fixed, and lock-step, with the same "skill-and-drill" lesson given to every child. Such an approach to teaching beginning reading conflicts with the basic principles of differentiated instruction, which are described in the next section. Second, most phonics programs incorrectly assume that children can only acquire knowledge of letter-sound patterns through direct instruction in which the teaching of letter-sound correspondences is explicit and systematic. The major difficulty with this assumption, however, is that there are simply too many letter-sound relationships in English orthography for children to acquire by direct instruction, probably several hundred.

Much, if not most, of what children learning to read in English come to know about the orthography is acquired through implicit learning. As the reading attempts of beginning readers who have letter knowledge, phonemic awareness, and a firm understanding of the alphabetic principle become more successful, the orthographic representations of more words become established in lexical memory from which additional spelling-sound relationships can be induced without explicit instruction. In contrast to letter-sound correspondences acquired by direct phonics instruction, which are largely context free (i.e., involve one-to-one correspondences between single letters or digraphs and single phonemes), letter-sound correspondences acquired by implicit learning are mostly context sensitive (i.e., depend on position-specific constraints or the presence of other letters). As children continue to develop in reading, they will begin making greater independent use of letter-sound information to identify unfamiliar words in text. Once this point is reached, the most effective way that children can achieve further progress in learning to read is through print exposure, as reading itself provides practice opportunities for building fluency and for facilitating implicit learning of additional letter-sound patterns.

Although children must rely increasingly on induction to acquire the spelling-sound relationships necessary for learning to read, explicit phonics instruction plays an important role in helping to "kick start" the process by which beginning readers acquire untaught spelling-sound relationships through implicit learning. Phonics instruction is therefore best thought of as a means to an end, not an end in itself. Because of the nature of English orthography, one of the main functions of phonics instruction is to provide beginning readers with a process for generating approximate phonological representations of unknown words that gets them close enough to the correct phonological form that, with context, the correct identification can be made. Children learn to use their knowledge of

spelling-to-sound relationships acquired through phonics instruction to produce partial decodings for unknown words, especially those containing irregular, polyphonic, or orthographically complex spelling patterns. These approximate phonological representations provide the basis for generating alternative pronunciations of target words until one is produced that matches a word in the child's lexical memory and makes sense in the context in which it appears. Additional spelling-sound relationships, especially context-sensitive patterns, can then be induced from the stored orthographic representations of words that have been correctly identified.

Phonics instruction is useful not because of the specific letter-sound correspondences taught (which are limited in number), but because it instills in beginning readers a firm grasp of the alphabetic principle, and gives them practice in looking closely at word spellings. Some explicit phonics instruction may therefore go a long way in facilitating the process by which children induce untaught spelling-sound relationships. For children encountering difficulty in developing the ability to perceive intuitively the redundant patterns and connections between speech and print, explicit instruction in alphabetic coding skills is crucial.

Teaching word analysis skills has been found to be more effective under some conditions than others. Providing beginning and struggling readers with explicit and systematic instruction in orthographic patterns and word identification strategies outside the context of reading connected text is more effective than only teaching word analysis skills incidentally (i.e., "as the need arises") during text reading, as is common in the whole language approach to teaching beginning reading. There are two reasons for this. First, instruction in word analysis skills that is deliberately separated from meaningful context allows children to pay full attention to the letter-sound patterns being taught, as well as avoid having their text reading overly disrupted. Second, isolated word study helps to ensure that beginning readers see the importance of focusing on word-level cues as the most useful source of information in identifying words, and to overcome any tendency they may have to rely primarily on sentence context cues in identifying unknown words rather than using context to supplement word-level information. One of the major distinguishing characteristics of struggling readers is their tendency to rely heavily on sentence context cues to compensate for their deficient alphabetic coding skills.

Although beginning readers should receive explicit instruction in letter-sound patterns outside the context of reading connected text, they should also be taught strategies on how and when to use this information during text reading through demonstration, modelling, direct explanation, and

guided practice. This includes teaching children to adopt a "set for diversity" in which partial decoding attempts are used to generate alternative pronunciations of target words until one is produced that matches a word in their spoken vocabulary and is appropriate to the sentence context. In general, better outcomes are achieved when phonics instruction is accompanied by rich and varied opportunities for children to practice and receive feedback on applying their newly acquired word analysis skills while actively engaged in the processes of reading and writing. Phonics instruction needs to be fully integrated within the literacy curriculum, not segregated from it.

Table 1 presents a continuum of teaching approaches for facilitating the development of alphabetic coding skill that ranges from the isolated skill-and-drill approach to the whole language, incidental learning approach. The approach that we recommend, the metacognitive strategy teaching approach, is based strongly on contemporary research and falls between the two extremes.

Table 1. Continuum of teaching approaches for facilitating the development of alphabetic coding skill.

Isolated Skill-and-Drill Approach	Metacognitive Strategy Teaching Approach	Whole Language Approach
• Atomistic view of reading acquisition; reading broken down into several subskills	• Dynamic view of reading acquisition; child seen as active learner	• Reading acquisition seen as natural process that is meaning driven; "no meaning, no gain"
• Heavy emphasis on teaching subskills in isolation; much seatwork and use of workbooks	• Emphasis on developing self-improving strategies for recognizing words and on how and when to use such strategies	• Minimal emphasis on word analysis activities; should only arise incidentally in context of reading connected text.

A criticism of the metacognitive strategy approach to teaching word recognition ability, however, is that English orthgraphy contains so many irregularities that focusing too much attention on teaching alphabetic coding skill not only wastes valuable time but possibly even confuses children and impedes progress. As noted previously, when examined at a deeper level, English orthography is not nearly as irregular as often claimed. Moreover, no word in English is completely phonologically opaque. Even irregularly spelled words like *stomach*, *castle*, *spinach*, and *friend* provide accurate phonological cues to the word's identity. Learning to read would certainly be a much more difficult task if spoken words like *stomach* were represented in the orthography as a random sequence of letters (e.g., *omtshca*) rather than as irregular spellings. When beginning readers apply their

developing knowledge of spelling-to-sound relationships to unfamiliar irregular words, the result will often be close enough to the correct phonological form that sentence context can be used to arrive at a correct identification, provided that the word is in the child's listening vocabulary.

Not only is alphabetic coding skill necessary for learning to read irregularly spelled words, but it is even necessary for taking advantage of the constraints of sentence context in identifying unfamiliar words in text. Only children who have begun to acquire alphabetic coding skill are able to benefit from sentence context. For children with limited alphabetic coding skill, context provides little or no help in identifying unknown words, as the words are largely unanalyzed. For example, beginning readers who are able to generate the regularized pronunciation "stow-match" when confronted with the word *stomach* are more likely to identify the word when it appears in the sentence context, "The football hit him in the *stow-match*," than children who fail to identify any of the sounds of the word, or perhaps only the sound of the first letter or boundary letters; for example, "The football hit him in the *statch*."

Those who claim that English orthography is too irregular to be of much use typically argue that children should be taught to use multiple cues in identifying words in text, with text-based cues (e.g., picture cues, sentence context cues, preceding passage content, prior knowledge activated by the text) being used to generate predictions about the text yet to be encountered and letter-sound information generally being used for confirmation and self-correction. As evidence against this claim, consider the results of the following study that examined the relationship between beginning readers' reported strategies for identifying unknown words in text and later reading achievement. Five-year-old (Year 1) beginning readers participating in a 3-year longitudinal study of early literacy development in New Zealand were divided into three groups according to their responses to the following question: "When you are reading on your own and come across a word that you don't know, what do you do to try to figure out what the word is?" The majority of children (52.1%) reported using word-based strategies, 34.3% reported using text-based strategies, and 13.6% of the children did not provide a response.

Typical examples of word-based strategies reported by the children were:

"Sound it out."
"Think of the sounds."
"Say the letters."
"Do the sounds of it."

"Make the sounds."

"Hear all the letters."

"Listen to what the letters are."

"You try and get the letters right."

"You hear the letters."

"Say out the sounds in the word."

"Sound it out, dad says so."

Typical examples of text-based strategies reported by the children were:

"Guess."

"Think, guess what the word is."

"Read it over again."

"Read on."

"Have a look at the picture."

"Keep on going, then go back and see what the word is."

"I leave it."

"Think about the word. Try to guess what it is."

"You just read it back again."

"Go back to the beginning, then you read it again."

"Miss it and go to the end and go back and guess a word that makes sense."

"Get your mouth ready."

These results indicate that the text-based word learning strategies emphasized by the classroom teachers in this study were not necessarily reflected in the word identification strategies that the majority of children actually used in learning to read. Of greater importance, the results further showed that the Year 1 beginning readers who reported using word-based strategies strongly outperformed the children who reported using text-based strategies on all reading and reading-related measures taken in the middle of year 3. Moreover, these children were six times less likely to enter Reading Recovery in Year 2 than the children who relied on text-based strategies in Year 1 (6% versus 37%). These findings provide further support for the view mentioned earlier that making use of letter-sound relationships to identify unfamiliar words in text is the basic mechanism for acquiring fast, accurate word recognition skill.

Assessment and Differentiated Teaching

Figure 10 presents the overall Cognitive Foundations framework combining the cognitive elements underpinning the development of the language comprehension (see Figure 5) and word recognition (see Figure 9) components of the Simple View of Reading. The structure of the framework is not intended to suggest that the development of the more advance cognitive elements cannot occur until all of the more fundamental elements are fully developed. Although some level of mastery of the more fundamental elements of the framework is needed to develop mastery of the more advanced ones, the elements tend to develop congruently and reinforce each other in a reciprocally facilitating manner, as noted previously (see figure 2). The elements of the framework should therefore not be taught in isolation from each other but instead should be taught in a more integrated manner; beginning readers should be given plenty of opportunities to practice and receive feedback on applying their newly acquired skills while engaged in performing the more advanced cognitive functions specified in the model.

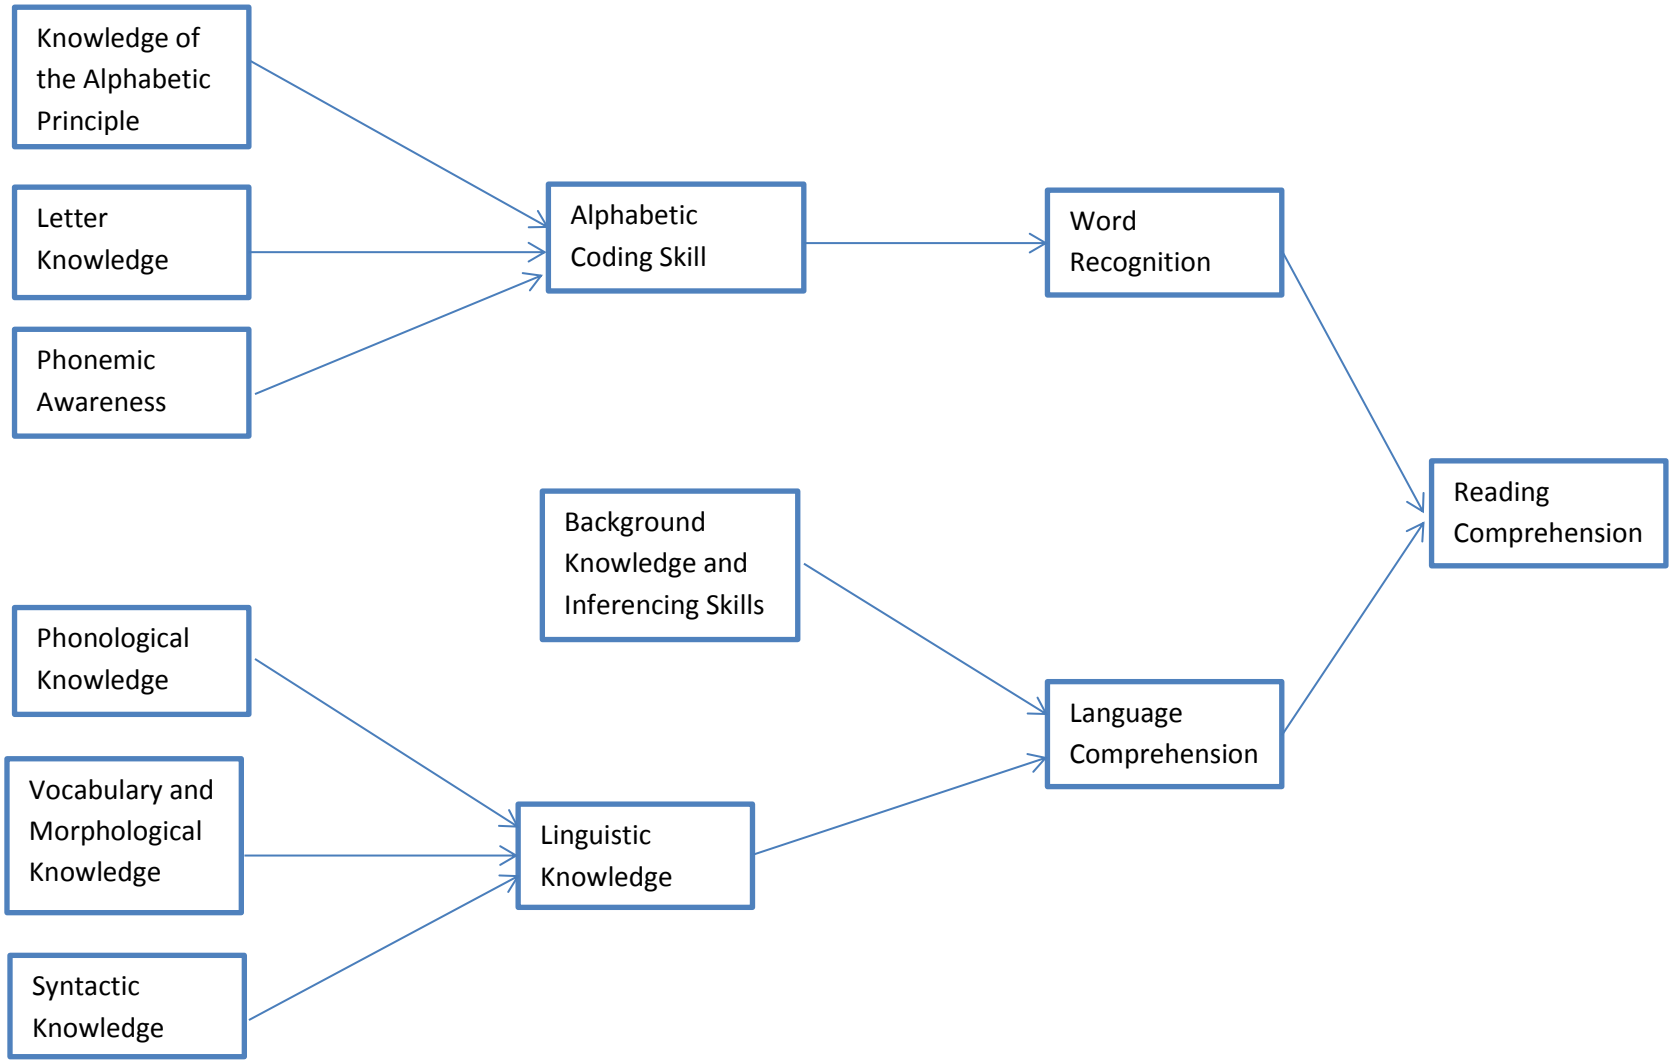

Figure 10. The Cognitive Foundations of Learning to Read

Underlying the Cognitive Foundations framework is the assumption that learning to read follows a developmental progression from pre-reader to skilled reader that involves qualitatively different but overlapping phases. Skill in comprehending written text depends on the ability to recognize the words of the text accurately and quickly; the development of automaticity in word recognition in turn depends on the ability to make use of letter-sound relationships in identifying unfamiliar words; and the ability to discover mappings between spelling patterns and sound patterns in turn depends on letter knowledge, phonemic awareness, and knowledge of the alphabetic principle. The literacy learning needs of beginning readers necessarily vary because they differ in the amount of reading-related knowledge, skills, and experiences they bring to the classroom, in the explicitness and intensity of instruction they require to learn skills and strategies for identifying words and comprehending text, and in their location along the developmental progression from pre-reader to skilled reader.

This points to the importance of differentiated teaching, where teachers use evidence-based assessment procedures and instructional strategies to cater to the different literacy learning needs of beginning readers from the outset of schooling. The structure of the Cognitive Foundations framework provides the basis for diagnostic reading assessment. For example, if beginning readers are not progressing satisfactorily in learning to derive meaning from print (i.e., reading comprehension, as shown in the far right box of the Cognitive Foundations framework presented in Figure 10), it is because they are having problems understanding the language being read (i.e., language comprehension), problems recognizing the words of text quickly and accurately (i.e., word recognition), or both. Weakness in word recognition skill stems from insufficient explicit instruction in alphabetic coding skill or inadequate opportunities to practice and receive feedback on applying alphabetic coding skills while actively engaged in reading. If alphabetic coding skills are still weak despite exposure to explicit instruction and practice, it is because of inadequate knowledge of the alphabetic principle, letter knowledge, or phonemic awareness.

A similar strategy can be followed in identifying possible weaknesses in aspects of oral language comprehension. For example, children with weaknesses in vocabulary and morphological knowledge or syntactic knowledge will have difficulty understanding sentences, and children with weaknesses in phonological knowledge may not be able to hear the difference between words with different meanings (e.g., *thin* and *fin*). Children who have problems in relating the meaning of each new

sentence in spoken discourse to the meanings of the sentences that preceded it by making inferences based on background knowledge will have difficulty understanding stories and passages.

In support of differentiated literacy instruction from the outset of schooling is research indicating that the amount of explicit instruction in word-level skills needed to initiate the process of inducing untaught letter-sound relationships through implicit learning varies considerably across children. Some beginning readers seem to grasp the idea after having had only a few spelling-sound correspondences explicitly taught to them, whereas other children require a fairly structured and teacher-supported introduction to reading.

Supporting an interaction between student characteristics and approach to teaching beginning reading are the results of a study by Carol Connor and colleagues examining the effects of different instructional emphases on children possessing varying amounts of reading-related skills (letter knowledge, oral vocabulary, knowledge of letter-sound correspondences) at school entry. They found that children who began first grade with below-average reading-related skills made larger reading gains in classrooms that provided greater amounts of teacher-managed, code-focused instruction throughout the year than in classrooms that provided greater amounts of child-managed, meaning-focused instruction (which included less instruction on word analysis skills and more reading of trade books and writing of text). In contrast, for children with higher reading-related skills at school entry, greater growth in reading was achieved in classrooms that provided lesser amounts of teacher-managed, code-focused instruction and greater amounts of child-managed, meaning-focused instruction.

Of particular importance was the finding that when student characteristics were appropriately matched with instructional approach, the improvement in end of year reading scores resulting from good fitting instructional patterns varied greatly between children with high and low levels of reading-related skills at the beginning of first grade. For high reading-related skills children, better fitting instructional patterns (i.e., child-managed, meaning focused instruction) resulted in about half a grade equivalent gain in end-of-year reading scores over less well matched instructional patterns (i.e., teacher-managed, code-focused instruction). However, for low reading-related skills children, better fitting instructional patterns (i.e., teacher-managed, code-focused instruction) resulted in a difference of more than *two full grade equivalents* in end of year reading scores compared with poorly fitting instructional patterns (i.e., child-managed, meaning focused instruction). These findings underscore the importance of implementing differentiated reading instruction from the

beginning of school, especially for those children with limited amounts of school entry reading-related skills.