Note. This is an author draft prior to final review. The official, final version will appear in ***The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, Fourth Edition*** **[J. Wixted, Ed.].**

Chapter Title: READING

Manuel Carreiras, BCBL.
Basque Center on Cognition, Brain, and Language
Ikerbasque. Basque Foundation for Science
Universidad del Pais Vasco (UPV/EHU)

Blair C. Armstrong
BCBL. Basque Center on Cognition, Brain, and Language
Department of Psychology and Centre for French & Linguistics at Scarborough, University of Toronto

Jon Andoni Duñabeitia
BCBL. Basque Center on Cognition, Brain, and Language

Abstract

Communication via written words is one of humanity's greatest inventions and plays a critical role in modern society. This chapter outlines the key cognitive, neural, and computational aspects of the reading system. In so doing, it shows how reading takes advantage of domain-general processing abilities and bootstraps written communication from other neurocomputational systems, including vision and spoken language processing. It also explains how failure in different parts of the reading system can lead to reading disorders such as dyslexia. Furthermore, emerging trends reveal exciting new directions for reading research, including advancing the understanding of how the brain changes as a function of learning to read, how the brain adapts to process different languages, and how to formalize our understanding of reading in more biologically plausible models. This chapter thus outlines how an interdisciplinary perspective to understanding reading has and will continue to advance our understanding of reading in ways that are critical for both fundamental and applied aims.

KEYWORDS: reading, visual word recognition, visual word form area, neural networks, dyslexia, neural plasticity, connectionist models

Introduction

Reading is a fundamental skill in modern society given the vast amount of information that is made available via the written word. Consequently, it is no surprise that literacy is a key determinant of a range of cognitive, social, and economic outcomes. Basic reading skills are now essential to our daily existence---for instance, to process information from signposts, text messages, address books, shopping lists, and medication instructions. Proficient reading skills also represent an invaluable component of more advanced socioeconomic opportunities, including postsecondary education, acquiring the latest scientific knowledge, and learning about who we are and where we came from based on our ancestors' own words. The societal changes brought about by the invention of reading are also not restricted to changes in our behavior alone. Paralleling these profound overt changes in the availability of information and in the time humans spend acquiring information through reading, researchers are also developing an increasingly deep appreciation for how these changes are impacting the representation and processing of information at a neural level (Carreiras, Seghier, et al., 2009), as well as of how reading abilities interact with and reshape other cognitive abilities (e.g., Behrmann & Plaut, 2013).

Collectively, unlocking the knowledge made available through the written word therefore makes reading one of the most exciting discoveries of our lifetimes---both from the perspective of individual readers, and from the perspective of scientists that study the details of the neural, cognitive, and computational underpinnings of reading. In the present chapter we will discuss (1) the cognitive representations processes involved in reading; (2) how the cognitive processes involved in skilled reading can be better understood by implementing computational models of reading and (3) what neural circuitry underlies reading abilities in the context of normal and impaired reading abilities, such as dyslexia. Finally, (4) we highlight some pressing unanswered questions and important future directions that can contribute to advancing our understanding of reading in the years to come.

Background Issues

I. The building blocks of the reading system

Reading is a parasitic ability that draws upon oral language processing and visual perception. In fact, in many ways reading can be conceptualized as a specialized type of visual perception akin to face recognition and other specialized visual object recognition abilities (Behrmann & Plaut, 2013), in that readers must extract a high level of detail regarding what line segments were presented where (i.e., there must be sensitivity to the configuration of the line segments) to discriminate between visually similar words that denote different meanings (e.g., pat, tap, bat, hat).

Perhaps the most salient aspect of how we process written words is, however, the fact that although reading shares some basic properties with general visual object recognition systems as well as specialized systems for some visual classes (e.g., faces), written words themselves have only existed for a few thousand years. This is an imperceptibly small amount of time in evolutionary terms, and contrasts starkly with the many millions of years during which relatively sophisticated visual processing systems have evolved, to say nothing of simpler visual abilities such as simple contrast or edge detection. As such, unlike these other visual abilities, at least some aspects of reading must be bootstrapped, borrowed, or recycled from other established cognitive systems. To preview our later section on neural circuitry, perhaps the clearest demonstration of the re-purposing of existing cognitive systems is in how components of the visual object recognition system are used by proficient readers to recognize visual words, which highlights just how early in the perceptual stream cognitive and neural optimization occurs in service of proficient reading. Additionally, the details of the brain lateralization of these optimizations of visual perception---which tend to manifest in the left hemisphere---are strikingly telling in terms of how other systems, such as those that subserve spoken language, are recruited to achieve proficient reading.

From a high level of abstraction, reading might be conceptualized as a relatively straightforward ability: individuals need to associate a particular set of lines that they see (e.g., the line segments that make up the visual word form D O G) onto the representation of a particular meaning (in this case, D O G denotes an animal that barks). However, this oversimplified description belies the rich and complex set of representations and processes that underlies the reading system. In service of moving towards a deeper understanding of reading, we therefore begin by unpacking the implications of this simple description as a starting point for identifying “where the action is at” in terms of relevant sub-components related to the reading process. A particular focus of this chapter is on how “indirect” mappings of a word’s spelling are able to activate the representation of a word’s meaning through the activation of a representation of the word’s spoken form. In the following sections we will briefly present a summary of the series of mechanisms that are, among others, involved in reading and that have been explored in depth in recent years. To illustrate the nature of some of the processes that lead to efficient reading, these mechanisms will be exemplified by some widely accepted empirical effects.

Orthography: Letter identity and position.

Years of research on orthographic processing have shown that when expert readers are presented with letter strings, they perform a series of automatic operations in order to access the lexical representation associated with those sets of letters, discriminating between potentially confusable neighboring orthographic representations. Nowadays, there is common agreement on that the basic-level orthographic operations needed to process strings such as the ones presented in Figure 1a can be accommodated within two main categories: those related to letter position coding, and those related to letter identity assignment. The former operations help the reader to disambiguate between words that share the same constituent letters, but in a different position (e.g., the words “tare”, “rate” and “tear”). The latter type of operations are critical in order to discriminate between words that share many of the letters in the same position, but only differ in the identity of minimal units (e.g., the words “rate”, “race” and “rake”), and more importantly, they are crucially involved in the identification of the same words written in different manners (e.g., written in upper case or lower case; “rate” and “RATE”).

Insert Figure 1 about here

In recent years, researchers in the visual word recognition domain have provided impressive and voluminous evidence from studies exploring letter position coding (Grainger, 2008). A clear exemplification of this work are studies investigating transposed-letter similarity (or confusability) effects (TL effects, hereafter; see Figure 1b). In a nutshell, the TL effect reflects the perceptual uncertainty with which readers initially encode the position of the letters that constitute a given string. When a reader is presented with a non-existing letter string (i.e., a nonword) that resembles an existing word but that only differs from it in the position of two of the internal letters (e.g., the nonword “cholocate”, which highly resembles the word “chocolate”), they consistently tend to take the visually presented nonword as the existing lexical representation (e.g., O’Connor & Forster, 1981; Perea & Lupker, 2003; Perea & Lupker, 2004; Perea, Palti, & Gomez, 2012; Schoonbaert & Grainger, 2004). In fact, readers take much longer to reject a nonword like “cholocate” than a nonword like “chotonate” in a lexical decision task, and they are prone to lexicalize items like “cholocate” and name them as “chocolate” in reading aloud experiments (Perea & Estévez, 2008). Furthermore, the TL effect has been replicated across many (but not all, as detailed later) languages in masked priming paradigms. In these tasks, readers perform a given task on consciously perceived target words (e.g., “chocolate”) that are very briefly preceded by transposed-letter nonword primes (e.g., “cholocate”) or by nonword primes in which the two critical letters are substituted by two other letters (i.e., replaced-or substituted-letter primes; e.g., “chotonate”).

As a consequence of the perceptual similarity between the masked and unconsciously perceived transposed-letter primes and the targets, responses to the latter are typically faster and more accurate in the TL condition than in the replaced-letter priming condition, thus resulting in a facilitative TL effect (Christianson, Johnson, & Rayner, 2005;

Dunabeitia, Perea, & Carreiras, 2007; Forster, Davis, Schoknecht, & Carter, 1987; Perea & Carreiras, 2006a, 2006b; Perea & Fraga, 2006, among many others). Interestingly, the TL effect survives extreme position manipulations, even though the magnitude of the effect is significantly greater for transpositions of contiguous letters than for transpositions of non-adjacent letters (Guerrera & Forster, 2008; Perea, Dunabeitia, & Carreiras, 2008b). Additionally, eye-tracking data have supported the assumption of an apparently effortless reading of sentences containing transposed-letter nonwords that highly resemble existing words as a result of a fast regularization process due to perceptual similarity (Rayner, White, Johnson, & Liversedge, 2006). The same authors recently showed that parafoveally previewing TL nonwords facilitated the reading of a target word as much as identical parafoveal previews (White, Johnson, Liversedge, & Rayner, 2008). Furthermore, electrophysiological correlates of TL nonword processing have replicated many of these findings, and shed light on the time-course of TL effects and letter position coding processes (Duñabeitia, Molinaro, Laka, Estévez, & Carreiras, 2009; Vergara-Martinez, Perea, Marin, & Carreiras, 2011).

Strikingly, these findings have been consistently replicated across languages, with the exception of Semitic languages. In those languages, reading is dramatically impeded when transpositions are for letters that belong to the roots of words (Perea, Mallouh, & Carreiras, 2010; Velan & Frost, 2007, 2009, 2011). This has motivated additional targeted studies that qualify the TL effect and its limitations even in non-Semitic languages (Dunabeitia et al., 2007; Dunabeitia, Perea, & Carreiras, 2009; Duñabeitia, Perea, & Carreiras, 2014; Lupker, Perea, & Davis, 2008; Rueckl & Rimzhim, 2011). Therefore, the bulk of evidence suggests that the TL effect is a solid marker of letter position assignment mechanisms, and ultimately, of basic-level perceptually mediated orthographic encoding.

What do all these TL effects tell us about letter position coding? The conclusions from all the studies reporting transposed-letter confusability or similarity effects seem to agree in a basic finding: early in the process of word recognition, the individual positions of the constituent letters are not coded in an accurate manner, but rather position assignment follow uncertainty and flexibility principles. This conclusion is at odds with models of visual word recognition that favor slot-coding strategies (i.e., each particular letter is coded in a given slot within the string in a position-specific fashion; see Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Grainger & Jacobs, 1996) and supports flexible position representations (see below and the modeling section for a detailed description).

The existence of TL effects brings a different critical question into play, which has been of interest to researchers exploring the importance of the letters and their location for decades: are all the letters we read really needed to get to the meaning? This question has been tackled from different angles in recent experimental research. Here, we focus on one of the most prominent and fruitful avenues used to explore the role of the letters and their positions, and their relative value within the string, via so-called relative position effects (RP effects, hereafter; see Figure 1c).

In brief, RP effects are typically found in masked priming research, and demonstrate that the relative internal ordering of the letters within a string is the most critical factor determining efficient lexical access, rather than the precise order in which the letters are presented. To illustrate this effect, take for instance the word BALCONY as a target. A RP prime for it would be the string BLCN, which is a partial or subset prime that respects the relative (but not absolute) ordering of the letters. When the effects of such a relative position prime are compared with the effects of a control unrelated prime like FTRM, it is typically found that the RP prime facilitates the recognition of the target word (Grainger, Granier, Farioli, Van Assche, & van Heuven, 2006; Grainger & Holcomb, 2009a; Peressotti & Grainger, 1999). Interestingly, these studies demonstrated that the preview benefits of a relative position prime are indistinguishable from those obtained from the preview of absolute position primes such as B_LC_N_, which respects the precise order of the letters within the string. This led scholars to propose that the way in which the order of the letters is processed necessitates a moderately flexible mechanism by which the relative position of letters is relied upon more heavily than the absolute and precise ordering of the letters.

Without undermining the fundamental existence of RP effects, it should, however be noted that, as it typically occurs with most psycholinguistic effects, the nature and extent of RP effects have been markedly qualified in recent years. For example, several authors have claimed that task demands can be responsible for the seemingly similar relative and absolute position priming effects reported in preceding studies. According to this account,under different experimental conditions that do not tax the system with rapid serial presentations, absolute position coding takes the lead in lexical access over relative position coding. Massol, Carreiras, and Duñabeitia (2016) followed on this line of reasoning in a series of explicit priming experiments. They demonstrated that two words sharing the same letters in the same precise position (i.e., absolute position) showed larger discrimination costs than two words sharing their letters in a different position (i.e., relative position). Furthermore, not all the letters seem to be equally important for RP coding strategies: Several studies have highlighted the preponderant role of consonants over vowels in this type of coding schemas based on flexible ordering (Carreiras, Dunabeitia, & Molinaro, 2009; Dunabeitia & Carreiras, 2011).

Over and above letter-position coding mechanisms, readers have to complete other visuo-orthographic stages in order to access a lexical or semantic representation. This is a major area of research, and the mental operations based on abstract letter identity identification that ultimately lead to efficient word recognition are not completely specified yet at every sub-stage of processing. However, we are gradually advancing towards a comprehensive theory of the architecture of the reading system, particularly using modern techniques withfine-grained temporal resolution. These types of studies highlighted a temporal sequence of mechanisms that guide abstract letter identification (Grainger, Rey, & Dufau, 2008). By investigating how readers access the mental representation of an abstract orthographic unit given a concrete piece of printed information in a particular color, size and form, a braod set of results support r different levels of letter identity assignment, in line with and adding additional detail to the levels of representation in classic models of visual word recognition (Jacobs & Grainger, 1991; McClelland & Rumelhart, 1981). For instance, the results of a masked priming letter identification ERP experiment showed that individual letter recognition follows (at least) three temporarily distinguishable stages (Petit, Midgley, Holcomb, & Grainger, 2006). In a first stage, readers extract sub-letter features (see also Mueller & Weidemann, 2012 for review on the influence of visual similarity in letter recognition). Next, readers access case-specific letter representations, and in a subsequent stage, readers access abstract-level information that are case-independent. Only after this last stage can a reader identify that the lowercase and uppercase versions of the same letter (e.g., “a” and “A”) refer to the same entity.

Case-independent visual word recognition is clearly an important marker of reading expertise. As stated above, readers have to develop a series of strategies to overcome the physical similarity that exists between different letters (e.g., “i” and “l”), but at the same time they have to remain relatively insensitive to differences based on the size and type of font in which a given letter is printed. More importantly, readers have to be able to access the internal orthographic representation of a letter regardless of the case in which this letter is written (i.e., uppercase vs. lowercase). This is achieved by accessing so-called abstract letter identities (ALIs, hereafter), which are case-independent orthographic representations (Arguin & Bub, 1995).

Whether or not ALIs are sensitive to physical similarity has been a topic of research and debate for the last 30 years (Proctor, 1981). One potential explanation for a number of apparently contradictory findings is that methodological differences in terms of the number of repetitions of each item within an experiment or of the task or paradigm used in a given study can generate strikingly different results (Bowers, Vigliocco, & Haan, 1998; Kinoshita & Kaplan, 2008; Ziegler, Ferrand, Jacobs, Rey, & Grainger, 2000). Nonetheless, it is now widely accepted that the initial stages of lexical access governed by visuo-orthographic processing are not dramatically affected by the presentation of letters in alternating case, that is, mixing uppercase and lowercase letters (i.e., the case alternation effect; see Figure 1d). Thus, the first stages of orthographic processing are governed by a case-independent type of processing of ALIs, and it has been recently demonstrated that case alternation does not hamper the access to the abstract lexical representations of the words during visual-word recognition (Perea, Vergara-Martínez, & Gomez, 2015; Reingold, Yang, & Rayner, 2010). Nonetheless, it has been shown via ERPs that the abstract representation of letters are preceded by a transient effect of physical similarity both in the Roman and in the Arabic alphabets (Carreiras, Perea, Gil-Lopez, Abu Mallouh, & Salillas, 2013).

The fact that access to lexical representations is not radically affected by the case in which the letters are presented clearly speaks to the great tolerance of the visuo-orthographic system to disruptions in the form in which a given letter (or letter sequence) is presented. This result may not be too surprising considering the vast amount of different fonts and handwriting styles that any reader has to face in everyday life. Clearly, proficient reading of such variable input requires a high degree of tolerance or flexibility in the identification of the precise glyphs or strokes and curves that constitute the letters to efficiently access the ALIs. In the extreme of this flexibility, recent research has highlighted two additional effects that illustrate just how tolerant to form the reading system can be: the mirror letter effect and the Leet effect (see Figures 1e and 1f, respectively).

Mirror invariance refers to the surprising capacity of the visual system to mentally rotate a given image so that both the image and its mirror reversal are treated as identical. Certain visual areas are equally responsive to mirror reversals of the same objects as a consequence of these mirror generalization principles (Gregory & McCloskey, 2010). Nonetheless, this interesting feature is a double-edged sword: Despite the convenience of this property for general visual object identification, it is evidently detrimental for reading, because words and letters have a canonical orientation. Given that mirror generalization is an intrinsic property of the visual system, will mirrored words be processed as if they were written in the canonical orientation during initial stages of visual-word recognition? Behavioral and electrophysiological data seem to suggest that this is precisely the case. Dunbar & MacLeod, (1984) showed that readers show Stroop incongruency effects even when words are printed backwards (i.e., a mirror-word Stroop effect). This finding was particularly informative because in the Stroop paradigm, access to the lexical representation is generally taken as an index of automaticity of reading. Moreover, Dunabeitia, Molinaro, and Carreiras (2011) showed that electrophysiological markers of masked priming effects are, at least during early stages of visuo-orthographic processing, identical for words that include mirrored letters and for words in which all the letters appear in the canonical orientation. This same finding has been replicated behaviorally by Perea, Moret-Tatay, and Panadero, (2011), who further demonstrated that masked behavioral priming effects for mirror-letters are found as long as the reversed letters in the primes have their own representation (i.e., avoiding the inclusion of non-reversible letters like “b” or “p”). Similar evidence was obtained in an eye-tracking study reported by Duñabeitia, Dimitropoulou, Estévez, and Carreiras (2013), who investigated the development of mirror-letter processing as a function of reading expertise and demonstrated that beginning readers showed greater tolerance to mirror-letters than expert readers.

Finally, it is worth mentioning a concrete instance of flexibility-based letter identity coding that has been focus of research and debate in recent years: the Leet effect (see Figure 1f). Leet is a non-conventional, alternative alphabet that is used mostly on Internet communication and that was initially conceived to overcome the boundaries imposed by spam-detection firewalls. The core assumption of Leet writing is precisely that the reading system is highly tolerant of form variations, and it builds on the fact that readers will access ALIs even if the letters are replaced by letter-like characters that preserve the original letters’ shape. Hence, according to the guiding principles of Leet writing, a word like CAT could be represented by C4T, taking for granted that readers would process the latter as the former with minimal difficulty. Strange as it may appear to be, experimental evidence has supported this initial intuition, demonstrating that the use of letter-like characters does not dramatically impact visual-word recognition: as as long as the characters used in the replacement keep to a great extent the original letters' core form, the flexible nature of the initial visuo-orthographic stages helps overcome the impact of these alterations and guides the reader to the access of the ALIs. In the seminal study by Carreiras, Dunabeitia, and Perea, (2007; see also Perea, Dunabeitia, & Carreiras, 2008a) the authors demonstrated that a string containing letter-like symbols or digits such as M4T3R14L activate the base word MATERIAL as effectively as the unaltered word would in masked priming conditions, when the Leet string is presented as a prime and unconsciously perceived by the reader (Duñabeitia, Perea, & Carreiras, 2009; Molinaro, Duñabeitia, Marìn-Gutièrrez, & Carreiras, 2010). Interestingly, these Leet effects have been shown to be specific for letters, and elements lacking this same level of abstract identities (i.e., characters lacking ALIs) cannot be replaced by similar-looking characters of a different nature to generate the same effects (Kinoshita, Robidoux, Guilbert, & Norris, 2015; Perea, Dunabeitia, Pollatsek, & Carreiras, 2009).

All in all, the effects reported in Figure 1 and described in this section of the chapter indicate that after an initial stage in which the basic visual features of the visual form are extracted, readers necessarily move through a series of representations related with visuo-orthographic processing of the input. These representations constitute the grounds on which visual-word recognition facet of the reading system is built. Efficient reading is based on the correct recognition and processing of individual printed words, which constitute the primary building blocks of language processing, but it also requires a relatively precise coding of these letters' positions within the string. Accessing the semantic knowledge related to a printed word is undeniably preceded by the correct recognition of the identities and positions of the individual letters that constitute that given string, so that word processing is ultimately a convolution of visuo-orthographic factors that determine to great extent the later access to phonological and morpho-semantic units, which we describe in detail next.

Phonology: homophones and syllables

Visuo-orthographic processes are not the only ones that take a leading role in lexical access. While it is true that these mechanisms are the entry gate to the print, it is equally true that other pre-lexical processes occur prior to accessing the meaning of the strings that are being read (Grainger & Holcomb, 2009b). Decades of research have highlighted the importance of phonological processes in reading (Carreiras, Perea, Vergara, & Pollatsek, 2009; Ferrand & Grainger, 1993; Grainger, Kiyonaga, & Holcomb, 2006; Holcomb & Grainger, 2006; Ziegler et al., 2000). In fact, experimental research has demonstrated that even when one manages to completely partial out the role and impact of orthographic units, access to phonological codes takes place at initial stages of visual-word recognition (Dimitropoulou, Duñabeitia, & Carreiras, 2011). A good example of the automatic activation of a phonological representation can be garnered from the effects of homophony. Homophones, and particularly homophones that are not homographs (e.g., a gale-force WIND vs. to WIND and turn), are words that share their pronunciation (namely, their phonological representation) but that do not share their spelling and meaning, such as the English words MAID and MADE. Dozens of studies have shown that two words with complete or extensive phonological overlap do activate each other and may compete for lexical selection during visual-word recognition, in spite of the reduced orthographic overlap they may have (e.g., Bowers et al., 1998; Ferrand & Grainger, 1992, 1993; Lukatela & Turvey, 1990). As Rastle and Brysbaert (2006) comprehensively reviewed, phonological effects are rather ubiquitous and the effects of homophonic representations extend to a variety of reading paradigms and techniques.

For example,, Ferrand & Grainger (1992, 1993) and Grainger & Ferrand (1996) showed that a pseudohomophone prime facilitated the recognition of the target compared to a control pseudoword using a masked priming task. In addition, they found that orthographic priming (nerc-NERF) produced a facilitative effect with a stimulus onset asynchrony (SOA) of 33 ms, but phonological priming (nair-NERF) started to emerge only at a 55 ms SOA. These effects suggest that phonology, as well as orthography, are early processes in the activation of a lexical entry in the cognitive system. Further related evidence was provided by Pollatsek, Perea, and Carreiras (2005). They manipulated the context dependent letter “c” in Spanish and showed that prime target pairs that involved a phonological change (cinal-CANAL, /z/-/k/) were recognized slower than pairs that did not imply such change (conal-CANAL, /k//k/), with respect to their orthographic controls (pinel-PANEL, ponel-PANEL) at 60 ms SOA. Additionally, Carreiras, Ferrand, Grainger, and Perea (2005) compared the difference in recognition times of words preceded by a phonologically similar prime by the first syllable (fomie-FAUCON) with respect to a substitution control (fémie-FAUCON), to words which were preceded by a phonologically similar prime by the second syllable (retôt-GATEU) with respect to a substitution control (retin-GATEAU). Phonological priming occurred only in the first case, implying that phonological processing might be sequential.

Orthographic and phonological effects observed in behavioral tasks also modulated ERP waveforms in early time-windows (Carreiras, Perea, et al., 2009; Grainger, Kiyonaga, et al., 2006). For example, Grainger, Kiyonaga, et al. (2006) used the masked priming procedure to test whether the activation of the representation of orthographic and phonological primes could be reflected on the ERP waveform. Subjects were presented with words preceded by transposed letter primes (barin-BRAIN) and two letter substitution controls (bosin-BRAIN), and by pseudohomophone primes (brane-BRAIN) and their controls (brant-BRAIN). They observed that both manipulations showed a negative component at 250 ms (N250), although orthographic priming generated a slightly earlier response (around 200-250 ms) than phonological priming (around 250-300 ms). Carreiras et al. (2009) also investigated the time course of orthographic and phonological effects in Spanish: the nonword “conal” primed CANAL (/k//k/) more than the nonword “cinal” (/z//k/) compared to pure orthographic controls (ponel-PANEL, pinel-PANEL). Phonological priming in the former case was observed in the 350-550 ms window, whereas orthographic priming in the latter case was observed in the 150-250 ms time window. This data support the behavioral evidence about the primary role of orthography followed by phonology in the activation of a word's representation.

Although the basic coding units are letters and phonemes, the mapping rules that determine how orthographic and phonological units are linked may entail combinations of letters (or letter clusters) and phonemes that, based on mapping regularities, give rise to sublexical structures such as syllables. Thus, efficient identification of words can be achieved by direct activation from graphemes (comprised of one or more letters) and phonemes to semantics or/and through the computation of sublexical units (e.g., syllables). Whether such sublexical units are created and relied upon during word identification, however, may depend on the consistency of phoneme-grapheme correspondences in different languages. One likely candidate for syllabic processing is Spanish, given the regularity graphemes to phonemes and the consistent syllable boundaries. In a seminal paper, Carreiras and collaborators (Carreiras, Alvarez, & De Vega, 1993) tested the role of the syllable as a sublexical unit in word recognition in Spanish, using the single presentation lexical decision task. They used words that began either with a high or a low frequency syllable. Words with a low frequency first syllable were identified faster than words with a high frequency first syllable. They reasoned that these effects were attributable to the following process: Words with high frequency syllables would initially trigger a larger number of lexical candidates, and/or neighbors of higher frequency than words with low frequency syllables (high frequency syllables are shared by many more words than low frequency syllables). Therefore, it would take longer to select the correct candidate in a larger neighborhood or/and if there is a higher frequency neighbor in the syllable neighborhood. They also reported that words with higher frequency syllabic neighbors were recognized slower than words with low frequency syllabic neighbors (Carreiras & Perea, 2002; Conrad, Carreiras, Tamm, & Jacobs, 2009; Perea & Carreiras, 1998). Furthermore, using the masked priming technique, Carreiras & Perea (2002) found that similar syllabic primes (alto-ALGA) inhibited not only the recognition of the target compared to control syllabic primes (esto-ALGA), but also that primes that shared the syllabic structure of the target (zo.ta-ZO.CO) produced facilitation with respect to primes that did not share it (ziel-ZO.CO). These results suggest that (at least in some languages) the syllable is an important sub-lexical unit that operates at a pre-lexical level and that the number of higher frequency syllabic neighbors has an inhibitory effect on word recognition, together with syllable frequency.

The time course of activation of syllables has also been investigated with ERPs. Recent evidence has shown that sublexical units such as syllables modulate the P200 component. Barber, Vergara, and Carreiras (2004; see also Hutzler et al., 2004) manipulated the word frequency and the syllable frequency in a lexical decision task while recording ERPs. They presented a set of high frequency words and a set of low frequency words with their corresponding nonwords. Half of the words in each set began with a high frequency or a low frequency syllable. As expected, word frequency effects produced less negative amplitudes in the N400. In contrast, syllable frequency produced the inverse effect in the N400 (an inhibitory effect), while it showed a facilitatory effect in the P200 window.

This P200 effect has been replicated in a subsequent experiment studying syllable compatibility (Carreiras, Vergara, & Barber, 2005). In that investigation, words were presented in two different colors, so that the boundaries of the colors matched or mismatched the boundaries between syllables. The syllable congruency effect influenced the P200 and also the N400 windows, reinforcing the idea of sublexical and lexical processing. In all these experiments syllables would be computed at the sublexical level and influence word selection at the lexical level. For instance, high frequency syllable words imply greater activation early in processing, because of the activation of many candidates that share this first syllable with the target (P200 effect). This implies a harder process of lexical selection later on the N400, because the prospective candidates have to be inhibited to finally identify the correct word form and meaning.

In sum, sub-lexical phonological effects such as syllabic effects may be critical to correctly understand the role of phonology in the early aspects of printed word identification (Carreiras, Dunabeitia, & Molinaro, 2012; Carreiras, Vergara, et al., 2005; Conrad et al., 2009; Dunabeitia, Cholin, Corral, Perea, & Carreiras, 2010) , although it should be acknowledged that many of these effects are modulated by the properties of the language under study. Thus, some of these effects may only be related to a subset of orthographies with very transparent mapping correspondences and with very well defined syllable boundaries. In other languages where syllabic boundaries are not marked, other units could be more relevant. More research is needed in other orthographies to evaluate the extent to which the computation of sublexical units in polysyllabic words is the rule or the exception (for related discussion, see Share, 2008; Frost, 2012).

Morphology

The initial access to written words is dominated by a series of fast-acting mechanisms aimed at enabling the rapid activation of orthographic and phonological codes. But orthographic and phonological processing are not the only sets of mechanisms playing an active role in effectively guiding the reader from sensory input to meaning (Nation, 2009). In the last two decades, the role of pre-lexical morphology has been underscored (Amenta & Crepaldi, 2012; Rastle & Davis, 2008). As such, it is now widely assumed that readers extract basic morphological regularities from the visual input prior to accessing word meaning. Psycholinguistic evidence has demonstrated that when a reader is faced with a polymorphemic string such as the English word “walker”, the visual-word processor automatically strips the affix “–er” off and the stem “walk” is also accessed in the lexicon (Taft & Forster, 1975). Curiously, this seemingly automatic sub-lexical decomposition of morphological units sometimes yields incorrect lexical activation of semantically unrelated but superficially morphologically related elements. For example, readers activate the word “corn” as if it were the stem of the word “corner” by means of automatically detecting and stripping off any chunk that may resemble an affix (Diependaele, Duñabeitia, Morris, & Keuleers, 2011; Feldman, O’Connor, & del Prado Martín, 2009; Rastle & Davis, 2008; Rastle, Davis, & New, 2004). This appears to be a by-product of automatic morphological segmentation that is useful in most cases, and the potentially detrimental or competing representations that may incorrectly be activated as a consequence of affix-stripping mechanisms are certainly limited, given that most languages base their derivational morphological systems in transparent constructions. Consequently, the initial stages of visual-word recognition leading to efficient lexical access, which is the core process of reading, are better understood as a convolution of visual, orthographic, phonological and morphological effects. These multiple factors conspire conspire to create a series of pre-lexical representations that ultimately give access to the meaning of a word by reassembling all the individual units and chunks that have been activated during these early stages mentioned above (see Nation, 2009).

Semantics

At its core, the aim of any communication, be it through a written, spoken, or other modality (e.g., sign language, Braille), is to map a sensory input onto a semantic representation which denotes the meaning of the word. Representing meaning, however, is a much less tangible issue than representing the orthographic or phonological properties of a word, which are much more closely tied to a sensory input. Interestingly, the anterior temporal lobes, which represent a particularly critical component of the semantic representation network, are positioned near the center of the brain, which makes them well suited for integrating information from projections originating in all sensorimotor systems. As such, at least some component of word meaning appears to be denoted by the interconnections or affordances that exist between an object’s representation in multiple dimensions (e.g., the “meaning” of a cup is in part denoted by the fact that it visually appears to be a container with a handle attached to it, and that such handles can be associated with a motor representation associated with grasping a handle; Patterson, Nestor, & Rogers, 2007). Additionally, word meaning may also be defined in part by the associations that exist between words themselves (e.g., the word DOCTOR has a meaning related to NURse because those two words often occur together; Tranel, Logan, Frank & Damasio, 1997).

In the context of written words, tapping onto meaning is thought to occur in two main ways in the context of proficient readers: either through direct mappings between a visuo-orthographic representation in semantics, or indirectly through a mediating representation in auditory/phonological brain areas (for a review, see Carreiras, Armstrong, et al., 2014). Which pathway is used for accessing the meaning of each individual word is a complex issue, however, and appears to depend on a number of issues such as reading proficiency, the presence of a reading disorder, and whether the word in question follows the spelling-sound regularities of a language or violates them (e.g., the “I” in MINT, HINT, TINT vs. in PINT; for discussion, see Armstrong, Dumay, Kim, & Pitt, accepted).

Additionally, and beyond the scope of a comprehensive review here, proficient reading draws upon a range of other cognitive abilities, including working memory systems (e.g., Christopher et al., 2012), inhibition/selection systems (e.g., to select the contextually appropriate meaning of an ambiguous word such as BANK, which can refer to a river or to a financial institution; for a review, see Armstrong & Plaut, 2016), representational binding (e.g., to form unified representations of compound words such as SMARTPHONE and AUTOMOBILE; for discussion and a mechanistic account, see Mareschal & French, in press), as well as general statistical learning mechanisms to extract the statistical structure of a language (Frost, Armstrong, Seigelman & Christiansen, 2016). In a sense then, the reading system leverages most of the major perceptual and cognitive systems available in the brain to enable efficient communication of information, and a full understanding of reading would require a full understanding of cognition more generally. In the interim however, how can we gain traction on the operation of the reading system and make sense of it all? In the following sections, we review how computational modeling, neuroimaging, and the combination of these two techniques may be combined to reconcile the rich and complex interactions between multiple representations and cognitive systems to enable proficient reading.

II. Computational models of reading

One particularly fruitful means of understanding the reading system is by building computational models of various aspects of reading, such as how line segments combine to make words to enable visual word recognition, or how written words map onto phonology to enable reading aloud. In so doing, researchers are able to benefit from the numerous strengths associated with building and running simulations of how the mind and brain subserve language. First and foremost, in contrast to simpler verbal descriptions of the reading process, building explicit computational models forces researchers to be explicit about factors such as: (a) what aspects of language information are represented as part of the reading process and (b) how processing mechanisms operate on those representations. As a concrete example, a researcher interested in the visual aspects of reading might posit that the visual system contains increasingly complex representations denoting the line segments in a word, individual letters, and finally, representations of the full word form. They would then need to specify the processing mechanism that allow for information about line segments to be combined together to identify individual letters. Additionally, models often include (c) learning mechanisms that specify how the model acquires particular representations, or learns to complete particular processes based on existing mechanisms (e.g., how to map a visual word form onto a spoken word form).

Specifying the characteristics of the computational system both in terms of representation and processing is valuable for a number of reasons. First, it forces the researcher to be explicit about a number of factors that might otherwise not be included in a verbal theory. This commits researchers to a particular theoretical interpretation that can readily be tested in a detailed quantifiable fashion, rather than in abstract qualitative terms. A model that successfully simulates a target pattern of behavior can therefore be thought of as, at minimum, an existence proof that a particular set of mechanisms could explain human reading behavior. In contrast, a model that fails to account for a particular pattern of reading behavior must not be a complete model of the human reading system. This need not imply that a model that does not capture every detail of a dataset is useless, however---specific types of failures of the model can nevertheless be revealing in terms of how the cognitive system does or does not operate, as has been the case in comparisons between parallel versus serial models of reading. Further, good models are often those that are the most parsimonious---in essence, those that explain the greatest range of data with the minimum amount of representational or processing complexity---because they can be reasoned about more readily and distill away from minor details to reveal the most critical aspects of the domain. For instance, a model of reading aloud that only includes monosyllabic words will miss out on the challenges that a model that does include multisyllabic words must overcome; however, the additional details needed to simulate reading multisyllabic words may cloud the interpretation of the model and the identification of what properties are critical to all words, be they multisyllabic or monosyllabic, and which additional mechanisms are critical for dealing with multisyllabic words specifically (for discussion, see Perry, Ziegler, & Zorzi, 2010; Plaut, McClelland, Seidenberg, & Patterson, 1996).

Second, when studying a phenomenon as rich and complex as reading, researchers inevitably approach the domain from a range of different perspectives, including different behavioral and neuroimaging techniques to probe language behavior and its neural underpinnings. Without denying the value of these investigations, the scope of any one empirical study is necessarily limited in terms of the variables that can be manipulated and the scope of issues that can be probed. Equally important, therefore, is developing a means of unifying the diversity of findings across disciplines, paradigms and methodologies, so that these techniques can mutually inform one another. This allows the strengths of individual studies to be combined to yield additional insight than that offered by any one study in isolation. Computational models are a powerful tool for such unification, because a single model can be used to simulate and understand data from a range of sources.

Third, and arguably most importantly, building models offers a critical test bed for exploring the implications of particular computational principles and for generating targeted novel predictions that can guide an empirical research agenda (McClelland, 2009). For example, a model aimed at simulating reading behavior might also be used to simulate the internal time-course of processing, thereby shedding light on the neural dynamics that give rise to that pattern of behavior (e.g., Armstrong & Plaut, 2016). This is particularly important as the complexity of the underlying computation increases and is shaded by many different variables (e.g., semantic richness, orthographic density, etc.), often in a nonlinear fashion, which can lead to strikingly different effects for different types of stimuli as a function of how a range of factors interact.

The Connectionist Modeling Framework

To facilitate model development, researchers often do not begin by building models ‘from scratch’ each time they wish to simulate a target set of phenomena. Rather, similar computational machinery is re-used across many different models. This helps make the models more parsimonious and easier to understand because knowledge of previous models can be re-used. Additionally, the successes and failures of re-using assumptions from other models can in and of itself provide insight into the domain generality of the underlying mechanisms.

In the context of reading, one especially influential framework for developing such models is the connectionist, or parallel distributed processing (PDP) framework. This framework has been used to simulate word recognition from early visual and auditory inputs (e.g., McClelland & Elman, 1986; McClelland & Rumelhart, 1981) through to semantic outputs and multi-word integration (e.g.,Frank, 2006; McClelland, St. John, & Taraban, 1989; Rohde, 2002). Connectionist models are instantiated by connecting together large pools of neuron-like processing units, with each pool typically denoting a distinct type of representation and associated neuroanatomy (e.g., separate pools of units could stand in for the orthographic vs. phonological systems, as well as the brain regions that map between these two systems). The activation of one such pool of units can then spread out to activate units in other representational pools---for instance, the activity generated in a pool of units denoting visual/orthographic processing when a visual word is presented can spread out to activate representations in the phonological and semantic systems. The amount of information flowing between units is governed by the strength of the connections that exists between them.

Critically, activation flows between different pools of units in a continuous fashion. This type of cascaded/interactive processing allows for two pools of partially activated units to constrain one another and help resolve coherent representations at multiple levels of representations simultaneously. This type of processing dynamic has proven critical for developing theories of phenomena such the word superiority effect---wherein an individual letter is recognized more rapidly in the context of a word than in the context of a wordlike nonword (Reicher, 1969)---using the Interactive Activation model (McClelland & Rumelhart, 1981; Figure 2), discussed in detail later. The inherent interactivity of connectionist models delineates them starkly from classic staged and modular theories and models (in the vein of Sternberg, 1969), in which processing at one level of representation (e.g., orthography) must be completed before engaging processing at a subsequent level of representation (e.g., semantics; for additional discussion, see Plaut & Booth, 2000; Borowsky & Besner, 2006; Armstrong & Plaut, 2016).

By virtue of their implementation, connectionist models have a number of important strengths. First, these models are said to be “domain general” because the same underlying computational principles can be used for understanding a range of cognitive abilities. Thus, models built with the connectionist framework are well suited for integrating across domains, such as by integrating a model of reading with a model of decision making to simulate reading performance in a particular task such as lexical decision. Second, by being grounded in an abstraction of the neural systems that operate to process information in the brain, these models are naturally suited to make contact with a range of data sources (e.g., neural data, behavioral data). Third, by being integrated with learning theory, models can explain not only proficient reading, but the learning trajectories that lead up to proficient reading (e.g., how morphological inflection is learned in the context of regular and exceptional English past tense; for a critical review, see Seidenberg & McClelland, 2014). Finally, although connectionist models are built by interconnecting sets of relatively simple processing units, these models often produce emergent behavior whose complexity and/or unexpectedness gives rise to additional understanding and explanatory power than might have been predicted (McClelland et al., 2010). For instance, these models can provide emergent accounts for how readers generalize knowledge from a few newly learned words to other new words (e.g., Armstrong, Dumay, Kim, & Pitt, Accepted), or how knowledge is impaired following brain damage (e.g., Joanisse & Seidenberg, 1999).

In the following section, we review how a few connectionist models have interacted fruitfully with other alternative theoretical accounts to advance our current understanding of various aspects of the reading system. To date, these models have focused on specific sub-components/sub-pathways in the reading system, such as how readers recognize printed words, how readers map print to speech, and how readers map print to meaning, so the review too focuses on how information passes from early sensory inputs up to the representation of meaning.

Visual Word Recognition: The interactive Activation Model

One of the most widely cited models in the reading literature is the Interactive Activation Model, originally proposed by McClelland & Rumelhart 91981). As one of the earlier connectionist models, it served to establish the validity of the abstract brain-style computational principles that was used to implement it. Testifying to the basic validity of how it has shaped thinking about visual word recognition, the Interactive Activation model remains highly influential in theories of visual word recognition to this day. The core contribution of this model was the notion of constraint satisfaction across partially resolved representations of individual line segments, letters, and words. This occurs in both a cooperative fashion across levels of representation (e.g., the units denoting the letters, T, R, and A, and P each sent excitation to multiple words that contained those letters, but all sent activation to the unit denoting the word TRAP), and in an inhibitory fashion within a level of representation (e.g., the word TRAP inhibited the words TRIP, TAKE, and TIME). To preview a future section, the lateral inhibition also appears to be particularly prescient of more sophisticated neural activity regulation mechanisms (e.g., Laszlo & Plaut, 2012; Laszlo & Armstrong, 2014; Armstrong & Plaut, 2016). Despite the relative simplicity of the IA model, it was shown to be incredibly successful at simulating a range of effects, including the word superiority effect, and neighborhood effects (e.g., enhanced processing of words with higher bigram frequencies than lower bigram frequencies) among many others. In so doing, this model fundamentally undermined the notion of a modular, staged account of perceptual-to-cognitive coding (although for a recent revisitation of the hypothesis that there are no top-down constraints on perception, see Firestone & Scholl, 2015), and questions the validity of serial processing models of reading (although, see, e.g., Whitney, 2001, for an alternative account involving serial processing).

Insert Figure 2 about here

Issues with the IA model. The core strengths of the interactive activation model notwithstanding, subsequent research has revealed a number of important issues with the original IA model. One fundamental limitation was that this model was developed prior to the introduction of general muti-layer supervised learning techniques such as backpropagation (Rumelhart, Hinton, & Williams, 1988). As a result, the connection strengths between units in this model were tailored by hand, an issue that is not tractable for simulating larger vocabularies, modeling language development, or for integrating much broader sets of representations (e.g., phonology, semantics). Subsequent models in this vein have all tended to involve learning mechanisms to address this issue, which greatly broadened the nature of the representations that are considered as part of the model (e.g., see Grainger & Ziegler, 2011, for discussion).

Another fundamental issue is the limited “slot” coding used in the IA model to code for letter inputs. In that model (and as occurs in an analogous fashion in related models of spelling-sound correspondence; McClelland & Elman, 1986), letter position is coded across specific slots, each corresponding to a fixed position within a word. All slots are also created equal in the model, leading the model to predict incorrectly that there should be equal performance for transposed letter effects effects regardless of the position of the letter within the word, other considerations (e.g., the frequency of that bigram) being equal. For example, the model fails to account for the asymmetric U-shaped distribution with which letter position and letter identity are detected across a range of tasks (Gomez, Ratcliff, & Perea, 2008).

To address this limitation of the IA model, subsequent connectionist and other models have adopted a number of schemes to make orthographic coding more flexible (see e.g., Davis, 2010; Gomez et al., 2008; Grainger & Ziegler, 2011; Whitney & Berndt, 1999).

These representations have often been combined with coarser/faster and slower/precise orthographic representations in an attempt to do better justice to the fine-grained time-course of word activation and to the representations needed for reading through phonology, which involves a precise encoding of letter position, and which varies as a function of literacy (Ziegler & Goswami, 2006). In turn, these models have continued to achieve improvements over the original IA model on a wide range of measures of reading ability. However, as discussed later, this success may be overstated in some cases due to over-tailoring of the models to account for data in particular languages such as English or relatively narrow language groups (e.g., Western European languages) at the expense of explaining performance in other languages.

Mapping Spelling to Sound: The triangle model

A separate line of connectionist models has been focused on the mapping between spelling, sound and meaning, with particular emphasis on the importance of the mapping between spelling to sound in the context of reading aloud. The initial motivation for developing these models was to account for how individuals are able to read nonwords and words that share the regularity of a language (e.g., the "I" in words MINT, HINT, or TINT, or in nonword GINT), while at the same time being able to explain accurate performance for exception words (e.g., the "I" in PINT). According to one account, at least two distinct processing routes were necessary to account for these data: one sublexical route focused on extracting the simple regularities underlying the language (so-called grapheme-phoneme correspondence, or GPC rules) and another lexical route which focused on learning representations of entire words (Coltheart et al., 2001). So-called regular words could then be read either by the lexical or sublexical route, whereas nonwords would be read primarily by the sublexical route, and exceptional words would be read by the lexical route, so as to override the regular pronunciation that would be generated by the sublexical route. The potential viability of this account has also been established by the model's ability to account not only for a range of data form proficient readers, but also from the model's ability to be subjected to simulated brain damage and recapitulate a range of dyslexic behaviors (see, e.g., Nickels, Biedermann, Coltheart, Saunders, & Tree, 2008).

Accounting for similar behaviors using the connectionist framework is challenged by the fact that unlike in dual route accounts, which involve two distinct pathways and associated qualitatively different processing mechanisms, the connectionist formalism uses the same simple neuron-like processing units to connect different pools of units. However, with the advent of connectionist learning mechanisms, it was discovered that connectionist models are, in fact, able to explain the bulk of the effects accounted for by the dual route account directly in the mappings between orthography and phonology by exploiting nonlinear processing dynamics (Plaut, McClelland, Seidenberg, & Karalyn Patterson, 1996; Seidenberg & McClelland, 1989). Recent work has further revealed that these nonlinear dynamics are able to overcome the paradox of representing regular and exception words in this single pathway through representational warping (Kim, Pitt, & Myung, 2013): whereas regular words are represented by units with intermediate levels of activation across large numbers of units, exceptional items force the units to extreme activation states (essentially completely on or completely off) which causes a localized distortion in the representational space where a different set of "rules" can apply. Models that have gone even further and included an indirect semantic pathway (i.e., the model has two anatomical "routes" but unlike the DRC, both of these "routes" are built using the same simple neural processors; Plaut, 1997) have shown additional improvements in quantitative performance by allowing orthographic, phonological, and semantic information to interact.

Recent developments and future directions.

Towards a "Universal" Model of orthographic representations. Models of word recognition have grown increasingly sophisticated over the past 30 years, showing great success at accounting for data from a range of tasks with ever-increasing precision. However, a central concern as these models move towards asymptotic performance is that they are becoming over-tailored to reproducing a limited set of target data, primarily from English (Share, 2008). As a result, there is increasing concern that these models are not general language models. Recent computational research, however, is beginning to change on this front, with an increasingly large focus on modeling linguistic diversity (e.g., Conrad Perry, Ziegler, & Zorzi, 2014; Yang, McCandliss, Shu, & Zevin, 2009).

Lerner, Armstrong, & Frost (2014) exemplify this paradigm shift in their comparisons of how a connectionist model, which used overlapping representations of letter position as inputs to the perceptual system, could learn differential sensitivity to letter position when trained on English versus Hebrew. As observed in several empirical tasks such as TL priming and rapid serial visual presentation, English words showed equal facilitation for transposition primes as for identity primes. In contrast, the model trained on Hebrew showed extremely detrimental effects of letter transposition, just as was observed in analogous behavioral experiments. Further investigations of the training corpus revealed that this insensitivity was not due to a general increase in orthographic density in Hebrew (which has fewer letters than in English), but was instead due to the preponderance of anagrams in that language. As such, the visual system in the context of Hebew must necessarily develop additional sensitivity to letter position, over and above that typically needed in English, particularly for extracting meaning directly from print. The structure of the language therefore interacts with the general model architecture to emphasize different statistics in different language contexts. From these principles, it can also be readily inferred that languages which have intermediate amounts of word confusability based on changes in letter position would be expected to fall somewhere in between English and Hebrew in terms of sensitivity to letter position.

This work highlights the importance of developing more “universal” language models that are sufficiently flexible so as to be able to interact with a range of linguistic environments to generate different effects in different languages (Frost, 2012). It also highlights the importance of modeling cross-linguistic differences explicitly, so as to better understand how particular psycholinguistic variables can drive the learning of different functional architectures in the context of different languages. For example, whereas the indirect “semantic” pathway is necessary for reading aloud English exception words such as PINT, does the near-absence of exceptions in transparent languages such as Spanish enable a different use of the indirect pathway in reading aloud, or language processing more generally?

Modeling internal time-courses as opposed to behavioral end-states. Although many connectionist models try to make contact with neural data to some degree the bulk of the extant models continue to focus on accounting for overt behavioral performance. This is a fundamental weakness, in that it prevents the models from fully leveraging the rich evidence regarding reading that is available from a range of neuroimaging techniques such as fMRI, EEG, and MEG. In so doing, these models are open to the possibility that although they are able to arrive at the same behavioral end-state observed in behavioural experiments, the means by which they do so is not reflective of actual neural processing.

To address this issue, a number of recent modelling projects have focused on expanding the classic connectionist formalism to model the neural correlates of reading and meaning acess, such as can be accessed via ERPs during passive reading of words (e.g., Laszlo & Federmeier, 2011). These projects have revealed that the classic formalism is unable to account for the time-course of neural processing, as indexed by the brain’s electrical activity while reading (Laszlo & Plaut, 2012). However, by increasing the biological plausibility of the model in a few important respects, such as separating excitatory and inhibitory processing and more accurately modeling the distribution of excitation and inhibition, the models are successful at recapitulating both the end-state behavior and the neural activity associated with the internal processing lead up to that behavior. Interestingly, this is surprisingly reminiscent of the separation of excitation and inhibition in the original IA model, although lateral excitation as well as inhibition is possible in these models.

Several additional behavioral findings that are not well captured by the classic connectionist formalism also manifest themselves in more biologically plausible models, as well, such as the time-course of ambiguous word comprehension (Armstrong & Plaut, 2016). Further additions of simple principles from neuroscience, such as modeling the fatigue dynamics of the post-synaptic potential (Laszlo & Armstrong, 2014), has continued to improve the models’ ability to account for a range of effects, such as simple word priming effects, in both temporal and spectral ERP data.

Collectively, this initial work points to the value of future work that evaluates the trade-off between model simplicity and parsimony on the one hand, and fidelity to existing neuroscience knowledge on the other. This will help establish an update connectionist framework that is both parsimonious, computationally efficient, relatively transparent to understand, and reflective of the core principles of neural processing necessary to explain particular facets of reading-related behavior and its neural correlates.

The intertwined nature of learning, representation, and generalization. Computational models in general, and connectionist models in particular, have often emphasized the importance of being explicit about how a model learns and represents knowledge. This is important because it allows researchers to understand how learning mechanisms enable a system to derive a representation that can subserve a particular objective (e.g., mapping between spelling and sound). Recent work by Armstrong et al., accepted has reinforced this claim and further argued that learning and representation are not simply separate facets of cognition, but are fundamentally intertwined. For example, when learning a new made up word that either follows the regularity of a language (e.g., GINT as in MINT, HINT) or that violates such a regularity (e.g., GINT as in PINT), a network must implicitly determine whether to leverage existing knowledge of a regularity when inserting a new word representation among existing representation. One option is to insert the new word leaving the representational space largely unchanged, which would mean extending the established regularity of the language. Another option would be to warp the representational space, so that a different "rule" can apply and an exceptional pronunciation can be stored. Words with more ambiguous pronunciations (e.g., MIVE, pronounced like GIVE/LIVE, not HIVE/DRIVE) would involve an intermediate amount of warping because they partially fit the regularities of the language.

Critically, this intertwining of learning and representation implies that generalization of new word knowledge will be critically determined by whether representational warping was needed to accommodate the new word or not. If no warping was needed, this implies that the new word embodied a widespread regularity and should be generalized readily. In contrast, if the pronunciation was exceptional, the restriction of the warped space to a small portion of the representational space (enforced by the other regular items in that general neighborhood) would impede the generalization of the new word's spelling-sound mapping. Words that follow more ambiguous rules would be warped to some intermediate degree and would be associated with an intermediate amount of generalization.

This discovery has critical implications for theories of language acquisition, because it indicates that the generalization of word tokens is not on a word-by-word level. Rather, it can be encouraged or discouraged to broad degrees by encouraging or discouraging the formation of warped representations (e.g., by presenting more than one item to demonstrate that multiple words exhibit a regularity; Apfelbaum, Hazeltine, & McMurray, 2013). This reshapes a classic debate in the computational modeling literature regarding whether regular and exception words are accommodated using a single set of computational principles (as in neural networks) or via qualitatively different processing mechanisms (as in dual route models). In particular, this new viewpoint stresses the importance of learning how representations of regular, ambiguous, and exception words are created via learning and generalized in novel contexts.

Insert Figure 3 about here

By being couched in a domain general learning, representation, and processing framework, the investigations of the warping mechanism also demonstrate the broader theoretical power of developing connectionist models of one domain such as reading: the same representational principles can be readily extended to other areas. For example, warping offers a natural explanation for why it is sometimes easier or harder to generalize the reading of new words in a second language, and how a first language spills over to impact words in a second language (Ijalba & Obler, 2015). The warping mechanism can also explain why generalizations of spelling sound correspondences in transparent languages such as Spanish or Serbo-Croatian are readily made---essentially no warping is required in those domains. At the other end of the spectrum, mappings between spelling and meaning are necessarily arbitrary in nature, so extreme warping is necessary to insulate representations from this arbitrariness and avoid incorrect generalizations (see e.g., Plaut, 1997; Plaut & Shallice, 1993). Even more broadly, warping may contribute to understanding why we have particular patterns of generalization in other quasiregular domains, such as learning to pronounce the past tense (for discussion, see Pinker & Ullman, 2002; Seidenberg & Plaut, 2014), as well as to our basic understanding of why regularities generalize in some statistical learning tasks but not others (Frost, Armstrong, Seigelman, & Christiansen, 2015). Studies of the basic mechanisms that underlie learning to read and to generalize existing language knowledge therefore attest to the power of domain general theories. They are naturally suited to extend the simulation of a single domain, such as how words are read aloud in English, to have a profound understanding for how it is that we allow different representations to co-exist and interact throughout the language system, and throughout other cognitive systems more generally.

III. The neural basis of reading

Having now outlined the core representations and processes involved in reading and how they could be instantiated in basic neurocomputational turns, we now turn to an in-depth treatment of the neural basis of our ability to read. Reading is clearly an instance of a late-acquired and complex cognitive skill that, as previously described, involves several elaborate representations and processing systems (e.g., orthography, phonology, morphology, semantics, etc.), which may be improved with training but also selectively affected by neurological injury or developmental disorders.

Despite a number of seemingly contradictory findings from neuroimaging experiments and the lesion-deficit correlation literature, there is convergence on: (a) a major role of the left hemisphere in reading; (b) a brain circuit involved in reading comprising a group of brain regions such as the posterior fusiform gyrus, the angular gyrus, supramarginal gyrus, posterior inferior temporal gyrus, mid temporal gyrus, inferior frontal gyrus, and premotor cortex. This complex reading neural network with temporoparietal (TP), occipitotemporal (OT), and inferior frontal (IFG), areas, mainly in the left hemisphere, seems to be functionally divided into two pathways: (1) a left temporo-occipital ventral stream involving the left ventral occipito-temporal cortex and the left ventral inferior frontal gyrus; that is, the left fusiform, middle and anterior temporal and the pars triangularis in the inferior frontal cortex, and (2) a left temporo-parietal dorsal stream comprising the superior temporal gyrus and supramarginal gyrus, the angular gyri and inferior parietal areas, premotor and the dorsal inferior frontal cortex, IFG pars opercularis (Carreiras, Mechelli, Estevez, & Price, 2007; Price, 2012; Pugh et al., 2000). Figure 4 shows and extensive activation of the reading network derived from contrasting lexical decision and naming words and pseudowords against corresponding control conditions (responding “no” in the lexical decision and saying “false” in the naming task to false fonts).

Insert Figure 4 about here

In much of the literature, it has been assumed that the ventral stream mainly supports the mapping from orthography to semantics, whereas the dorsal stream would underpin the indirect mapping between the orthography and semantics through phonology. In the ventral stream, the left fusiform gyrus (the so called visual word form area) has been considered a core region for visual word processing, and in particular for processing orthographic information. This region exhibits higher activation to words or pseudowords than to false fonts or checkerboards (Baker et al., 2007; Binder, Medler, Westbury, Liebenthal, & Buchanan, 2006; Carreiras, Monahan, Lizarazu, Duñabeitia, & Molinaro, 2015; Cohen et al., 2002; Dehaene & Cohen, 2011). It is active regardless of the spatial location of the stimuli (RVF or LVF), insensitive to typographic case and displays quite reliable localization across subjects (Cohen & Dehaene, 2000; Cohen & Dehaene, 2004; Cohen, Dehaene, Chochon, Lehericy, & Naccache, 2000; Cohen, Dehaene, Naccache, et al., 2000; Cohen, Henry, et al., 2004; Cohen et al., 2002; Cohen, Lehericy, et al., 2004). Thus, this area has been conceptualized as an important “hub” within the distributed network underlying printed word processing.

However, the functional role of the visual word form area role is still the subject of an important debate. It has been speculated that this area corresponds to the storage site of known orthographic forms and is involved in bottom-up pre-lexical orthographic processing (S. Dehaene & Cohen, 2011; Dehaene, Cohen, Sigman, & Vinckier, 2005; Dehaene & Changeux, 2005). An alternative view is that this brain area is involved in integrating visuospatial features from sensory inputs with higher-level associations, via both bottom-up and top-down connections (Carreiras, Armstrong, Perea, & Frost, 2014; Carreiras, Quiñones, Hernández-Cabrera, & Duñabeitia, 2015; Price & Devlin, 2011).

In a related debate, it has been proposed that phonological computations in the dorsal pathway (Jobard, Crivello, & Tzourio-Mazoyer, 2003). However, recent data suggest that computation in regions of the dorsal pathway are also associated with a coarse-grained orthography-to-semantics mapping, highlighting the role of parietal regions in early stages of visual word recognition (Carreiras, Monahan, et al., 2015; Carreiras, Quiñones, Hernández-Cabrera, et al., 2015; Carreiras, Quiñones, Mancini, Hernández-Cabrera, & Barber, 2015; Reilhac, Peyrin, Demonet, & Valdois, 2013). In particular, activation of superior parietal regions has been associated to tasks that require multi-element processing, such as the visual attention span task (Lobier, Peyrin, Le Bas, & Valdois, 2012; Peyrin et al., 2012). In addition, Carreiras, Quiñones, Hernández-Cabrera, et al. (2015) found evidence from perceptual tasks that involve letter identity and letter position may be involved in the earlier stages of visual word processing.

These results question the division of labour and interactions between the two main neural networks, according to which the dorsal route would be mainly involved in orthography to semantics through phonology and the ventral route in the direct mapping between orthography and semantics. This debate thus parallels debates regarding the division of labour in computational models discussed in the previous section, including the ability to access semantics indirectly via phonology and the need (or lack thereof) for qualitatively different routes for lexical and sublexical knowledge. In agreement with the basic premise of connectionist models, the current neural data is more consistent with the idea that the dorsal and ventral pathways cooperate during visual word recognition processes (Rosazza, Cai, Minati, Paulignan, & Nazir, 2009). In further support of this position, structural connectivity between regions of the two pathways (the posterior parietal cortex and the inferior temporal cortex) has been documented (Thiebaut de Schotten, Cohen, Amemiya, Braga, & Dehaene, 2012), which could also contribute to these interactions. Further studies will shed light on how the two pathways interact to give rise to visual word recognition, since the division of labour between the ventral and the dorsal routes seems to be an excessively simplistic hypothesis and likely to be modulated by cross-linguistic differences.

How “universal” is the brain reading circuit?

Reading in transparent orthographies with consistent letter-to-sound correspondences (e.g., Spanish or Italian) is thought to rely more on phonological processes, whereas reading in opaque orthographies (e.g., English) is assumed to rely more on lexico-semantic processes. These fundamental differences in the orthographic structure of written alphabetic systems and in the taxing on different cognitive processes have been linked to functional variations in the underlying brain circuit for reading. Thus, it has been hypothesized that orthographic depth might modulate the engagement of regions in the brain circuit along the dorsal and ventral pathways. Some cross-language research has revealed differences in the reading circuitry of readers of different languages. For instance, Paulesu et al. (2000) reported that although both Italian (shallow orthography) and English (deep orthography) readers recruited left-lateralized areas of the dorsal and ventral reading pathways, readers of the shallow orthography more strongly activated areas of the dorsal reading pathway compared to their peers, and readers of the deep orthography showed the opposite pattern (i.e., more activation in the ventral pathway). Similar results were reported by Das and collaborators (Das, Padakannaya, Pugh, & Singh, 2011) with English-Hindi bilinguals. Activation in the ventral pathway (left inferior temporal gyrus) was stronger when they were reading in English, and in the dorsal pathway (left inferior parietal areas) when reading in Hindi. Finally, Oliver, Carreiras, & Paz-Alonso (2016) reported a functional co-activation of the left vOT cortex and (i) regions of the ventral pathway for Spanish-English bilinguals reading in English (L2 deep orthography), and (ii) regions along the dorsal pathway in Spanish-Basque bilinguals when they were reading in Basque (L2 shallow orthography), suggesting a complex role for the division of labor between different brain regions between languages within individuals.

Collectively, these studies of the reading circuit clearly show that it is sensitive to cross-linguistic differences driven by orthographic depth factors. However, these studies focused more on the differences than in the similarities of the universal reading circuit between languages. Recently, Rueckl et al. (2015) showed that similarities in the activation of the reading circuit across very different languages (Chinese, English, Hebrew and Spanish) are more prominent than the differences (the fact that there were some differences notwithstanding). They obtained this important result by recruiting skilled adult readers of these four highly contrasting languages to perform an identical semantic categorization task to spoken and written words. Speech-print convergence emerged in a common brain circuit across languages, regardless of whether their writing system was alphabetic or logographic, opaque or transparent. There were, however, small differences in the *degree* of convergence in some regions. Specifically, speech-print convergence was higher for (transparent) Spanish than (opaque) English and Hebrew in the left SMG and the left SMA, whereas it was higher for English and Hebrew relative to Spanish in the left angular gyrus and in several ventral left hemisphere regions including the fusiform gyrus, MTG and ITG, along with right STG and MTG. Potentially reflecting the tailored functional division of labor that develops in the brains of speakers of different languages, areas with higher speech-print coupling for a transparent orthography were related with phonological processing whereas those higher for opaque orthographies were related to semantic processing. This work therefore stresses the importance of both “lumping” and “splitting” approaches to understanding the reading circuit. There are clear similarities between languages that can be lumped together to fruitfully understand what appear to be universal principles of the neural basis of language. However, these principles split off in the details in many cases to modulate how the fine-grained operation of the reading circuit.

The temporal dynamics of the reading circuit

Multiple neural resources must work quickly and in synchrony to automatically access the meaning of the arbitrary visual symbols of written languages. Linking neural operations with reading requires monitoring not only brain activation provided by hemodynamic methods such as fMRI, but also brain signals with a sensitivity that matches the speed of processing. EEG and MEG monitor the electric and magnetic fluctuations associated with neural activity with millisecond resolution. MEG also allows the sources of the measured signals to be localized with relatively high accuracy, unlike standard EEG (Hämäläinen, Hari, Ilmoniemi, Knuutila, & Lounasmaa, 1993). Thus, EEG and MEG are thus ideally positioned to address questions about reading, which is fundamentally characterized by its extremely rapid temporal signature.

A robust reading-related N170 ERP component (also reported as P/N150) has been detected in both EEG and MEG for contrasts between words versus other low-level visual control stimuli such as strings of meaningless symbols, forms, shapes, dots, etc. (Bentin, Mouchetant-Rostaing, Giard, Echallier, & Pernier, 1999; Brem et al., 2010; Carreiras, Quiñones, Mancini, et al., 2015; Eulitz et al., 2000; Maurer, Brandeis, & McCandliss, 2005; Maurer, Brem, Bucher, & Brandeis, 2005; Maurer, Han, & McCandliss, 2005; Tarkiainen, Helenius, Hansen, Cornelissen, & Salmelin, 1999). Such reading-related activity has been assumed to reflect an automatic, specialized process, as it has been detected in passive tasks that do not require decisions on the words (Bentin et al., 1999; Brem et al., 2010; Eulitz et al., 2000; Maurer et al., 2005; Maurer et al., 2005; Maurer et al., 2005; Tarkiainen et al., 1999). Similar effects reported as N/P150 have been found for orthographic processing using masked priming paradigms (Carreiras, Dunabeitia, et al., 2009; Carreiras, Gillon-Dowens, Vergara, & Perea, 2009; Holcomb & Grainger, 2006). Thus, according to EEG and MEG evidence, visual stimuli can be identified as letter-strings by about 150 ms after stimulus presentation, with left-lateralized activation in the inferior occipitotemporal cortex (Carreiras, Monahan, et al., 2015; Carreiras, Quiñones, Mancini, et al., 2015; Tarkiainen, Cornelissen, & Salmelin, 2002). This left-lateralized N170 (N/P 150) was linked to the similarly left-lateralized hemodynamic activation during visual word recognition (Brem et al., 2010) to support the idea that the left fusiform is a mandatory hub for early orthography processing. However, early activation for words and for words and pseudowords compared to symbols was also found at 120-130 ms in the temporal-parietal cortex (angular gyrus and intra-parietal sulcus) and at 200 ms in left inferior frontal gyrus and left temporal areas (Carreiras, Monahan, et al., 2015;), which stresses the question of exactly what type of information and representational constraint is being captured by these neural correlates of the time-course of processing.

As noted in brief earlier in this section, there are two main theoretical positions regarding the temporal dynamics of the reading circuit. One position considers that the flow of information is mainly bottom-up with the left fusiform being a mandatory early orthographic hub. According to this account, the left fusiform would be in charge of computing abstract prelexical orthographic representations (Dehaene & Cohen, 2011). In contrast, the second position suggests that orthographic representations are also computed by other brain areas in the dorsal pathway and that the left fusiform is a multimodal hub. This multimodal hub computes not only orthographic representations, but other representations as well, and receives not only bottom-up, but is also modulated by top-down information (Carreiras et al., 2014; Price & Devlin, 2011). Studies that have attempted to tease these possibilities apart have revealed an intricate and complex pattern of effects, which we take to support the second position most strongly. For example, it has been shown that the left fusiform is sensitive to lexical frequency (Kronbichler et al., 2004) and is activated not only by masked orthographic information but also for masked pictures, what implies the access to semantic processing (Kherif, Josse, & Price, 2011). In addition, top-down lexical influences have been found in the N250 component using the masked priming paradigm (Carreiras, Dunabeitia, et al., 2009; Carreiras, Gillon-Dowens, et al., 2009; Duñabeitia, Molinaro, et al., 2009). In particular, Duñabeitia, Molinaro, et al. (2009) found differential effects in the N250 in transposed-letter priming for word-word pairs (e.g. ‘casual–causal’) and for nonword-word pairs (e.g. ‘barin–brain’), suggesting that high-order lexico-semantic information constrains orthographic form-level processing in the N250. Additionally, Carreiras, Dunabeitia, et al. (2009) showed that masked subset priming of consonants (e.g., mln-melon) and masked full identity priming (e.g., melon-melon) did not significantly differ from each other already in the N250 component, whereas masked vowel subset priming (e.g., aio-amigo) and masked identity priming (e.g., amigo-amigo) did. These effects demonstrate that top-down lexical information modulates the N250 component given that consonants are more lexically constraining (in fact, similarly to the full word), than vowels in predicting word identity (Dunabeitia & Carreiras, 2011).

Other data consistent with a top-down flow of information during visual word recognition come from MEG and intracranial recording (Cornelissen et al., 2009; Thesen et al., 2012; Wheat, Cornelissen, Frost, & Hansen, 2010; Woodhead et al., 2012). Using different paradigms they show very early effects in anterior regions (100-200 ms) that suggest an interactive flow of information between frontal (e.g., the inferior frontal gyrus) and posterior (left fusiform) areas. For example, Woodhead et al. (2012) contrasted words to false fonts and used dynamic causal modeling top show feedback from the inferior frontal gyrus onto the left ventral occipitotemporal cortex within the first 200 ms provided the best fit for the data relative to a model with only feedforward connectivity.

Taken together, these studies of the time-course of processing, facilitate by measures with high temporal sensitivity in addition to spatial sensitivity, suggest an early activation of the inferior frontal region that sends higher-level information to constrain the representations computed in the left fusiform. We are still clearly far from understanding exactly how the information flows in the reading circuit and what sort of computations are done by different areas, as revealed by the extensive body of literature reviewed in this section. However, the current set of data clearly challenge the notion of exclusively bottom-up flow of information with the left fusiform as an early mandatory hub that computes orthographic processing, and are more consistent with interactive processing dynamics.

Dyslexia

Learning to read entails a substantial reorganization of the brain (Carreiras, Seghier, et al., 2009). Interestingly, whereas the majority of children learn to read without too much difficulty with appropriate training, some of them struggle with learning to read. These atypical difficulties in developing reading skills are not only manifest at the level of behavior, either. The activity generated in the brain's reading circuit systems during reading, as well as the functional and structural connectivity between the areas of the circuit is different for children with reading disabilities as compared to typically developing children specially posterior areas (Brunswick, McCrory, Price, Frith, & Frith, 1999; Paulesu et al., 1996; Rumsey et al., 1997; Rumsey et al., 1992; Rumsey et al., 1997; Shaywitz et al., 1998; for reviews see Pugh et al., 2010; Richlan, Kronbichler, & Wimmer, 2011).

For example, as a functional level, there is evidence that whereas skilled readers recruit and tune the reading circuit, developmental dyslexia is associated with a failure to recruit the occipito-temporal cortex (Richlan et al., 2011). In contrast, recruitment in the homologous regions in the right-hemisphere has been reported in some studies of adults with dyslexia (Pugh et al., 2000; Rumsey et al., 1999). More recent findings are also broadening the reported differences between typical and dyslexic readers, for instance, by indicating that there may be a disconnection between access to phonological representations in the inferior frontal gyrus (Boets et al., 2013).

Similarly, at the structural level, within region (grey matter volume) and between region (white matter tracks as indicated by fractional anisotropy (FA)) anomalies have also been associated with reading disabilities across these regions (Hoeft et al., 2007; Klingberg et al., 2000; Niogi, Mukherjee, & McCandliss, 2007; Richlan, Kronbichler, & Wimmer, 2012) and with the lack of exposure to reading (Carreiras, Seghier, et al., 2009). Additionally, during recent years the important role of subcortical structures and their interactions with the distributed cortical regions has become increasingly clear. For example, a recent study found activation anomalies in the putamen, caudate nucleus, thalamus, and cerebellum in poor readers (Preston et al., 2010). These findings indicate that increased attention to known brain pathways connecting cortical and subcortical regions is warranted to paint the full picture of the neural basis of reading disabilities.

Fully leveraging the insights from functional and structural studies of dyslexia to achieve such theoretical advance is, however, challenged by a number of issues. One theoretically salient issue is that it is hard to tease apart whether observed differences are the cause or consequence of dyslexia. For instance, neuroanatomical alterations in auditory regions have been documented in dyslexic readers (e.g., Ramus, 2003; Ramus, Pidgeon, & Frith, 2003; Ramus, Rosen, et al., 2003), but whether these deficits are linked to auditory sampling in general, reading in particular, and/or an interaction between reading and these neural systems has yet to be fully answered. Another important theoretical issue concerns what representations are being processed exactly and how they are being manipulated and are being reflected by a particular neural correlate. To probe these issues requires very careful experimentation coupled with detailed theories of dyslexia. For example, several researchers have recently hypothesized that there is a causal role for oscillatory activity in synchronizing or entraining the spectral properties of auditory stimuli at distinct frequency bands, and that these deficits are accompanied by specific hemispheric lateralization patterns (Giraud & Ramus, 2013; U. Goswami, 2011), thereby offering an important linking function between neural correlates and observed behavior.

This possibility has received support from careful of atypical brain synchronization at both syllabic (slow) and phonemic (fast) rates in dyslexics, together with an atypical hemispheric lateralization of neural synchronization and an atypical hemispheric asymmetry in cortical thinning (Lizarazu et al., 2015). As part of the same set of studies, it was also shown that the neural entrainment to the multiple frequencies of the speech signal and the differential patterns of causal connectivity across the brain network are all implicated in auditory sentence processing. For example, there was reduced functional connectivity between primary auditory regions and the left inferior frontal gyrus (a high-level phonological brain hub) in dyslexics, suggesting that during speech comprehension, this “dysconnection” hinders the feedforward communication from the right auditory cortex to LIFG in dyslexic readers (Molinaro, Lizarazu, Lallier, Bourguignon, & Carreiras, 2016). Thus, it seems that improper low-frequency acoustic entrainment affects phonological processing and, in turn, normal reading acquisition. This represents only one exciting new hypothesis that seeks to advance our understanding of typical and atypical reading, but gives a flavor for the type of integrated theoretical and neuroimaging approach, as well as the consideration of neural temporal dynamics that we expect will play an increasingly important role in reading theories going forward.

Conclusion

The core theories and findings outlined in this chapter, while no means comprehensive, provide a snapshot of the rich and complex representations and processing mechanisms that underlie reading abilities, and how these processes can go astray in various ways and lead to reading disorders. This review also reveals the critical value of adopting an interdisciplinary approach to understanding reading, which takes into consideration both the statistical properties of different languages, behavioral outcomes related to reading, the neural correlates of reading expertise, and how the reading system is created by borrowing from other established cognitive systems. These diverse backdrops reveal the intricacies of the reading system, but also pose a challenge unto themselves given their combined complexity. However, this complexity may be addressed through the development of neurocomputational theories that can simulate the rich interactions among the different components of the reading system and make theoretical claims explicit and quantifiable. More broadly, by considering the reading system from the perspective of a domain general learning, representation and processing framework we can gain insight into not only the reading system, but other related cognitive systems as well.

Although much is known about reading from the large body of prior work on the subject, there are clearly many questions left to answer to arrive at a comprehensive theory of reading, as well. The plethora of findings from diverse theoretical and methodological perspectives offers valuable guidance and targeted research for future investigations: Can learning to read lead to the development of what are effectively domain-specific abilities unique to reading, or does the reading system always maintain universal properties common to the cognitive systems that it bootstraps from such as visual expertise, spoken language, and multi-modal integration? To what degree does reading per se share processes with other related cognitive systems, such as numerical cognition, reasoning, logic, and the extraction of meaning from visual forms? What ramifications does the restructuring of the brain to enable proficient reading have for other cognitive systems and their relative performance, such as the ability to perceive other classes that require fine-grained discriminations, such as face processing, or recognizing specific types of birds, cars, or houses? To which extent letters are special as compared with other cultural stimuli like digits or symbols? Is there any critical period for learning to read? Do we employ the same mechanisms for learning to read in the childhood and adulthood? How do different ways of representing orthography like different types of alphabets (e.g. Roman and Arabic alphabets) or logographic symbols modulate our cognitive processes and brain mechanisms? Will be able to predict reading disabilities by studying the preliterate brain? These are but a small sampling of intriguing questions that are emerging as key targets for future work as a result of research to date into humanity’s greatest invention.

References

Amenta, S., & Crepaldi, D. (2012). Morphological Processing as We Know It: An Analytical Review of Morphological Effects in Visual Word Identification. *Front Psychol, 3*, 232.

Apfelbaum, K. S., Hazeltine, E., & McMurray, B. (2013). Statistical learning in reading: Variability in irrelevant letters helps children learn phonics skills. *Dev Psychol, 49*(7), 1348.

Arguin, M., & Bub, D. (1995). Priming and response selection processes in letter classification and identification tasks. *Journal of Experimental Psychology: Human Perception and Performance, 21*(5), 1199.

Armstrong, B. C., Dumay, N., Kim, W., & Pitt, M. A. (Accepted). Generalization from newly learned words reveals structural properties of the human reading system. *Journal of Experimental Psychology: General.*

Armstrong, B. C., & Plaut, D. C. (2016). Disparate semantic ambiguity effects from semantic processing dynamics rather than qualitative task differences. *Language, Cognition and Neuroscience*, 1-27.

Baker, C. I., Liu, J., Wald, L. L., Kwong, K. K., Benner, T., & Kanwisher, N. (2007). Visual word processing and experiential origins of functional selectivity in human extrastriate cortex. *Proc Natl Acad Sci U S A, 104*(21), 9087-9092. doi: 10.1073/pnas.0703300104

Barber, H., Vergara, M., & Carreiras, M. (2004). Syllable-frequency effects in visual word recognition: evidence from ERPs. *NeuroReport, 15*(3), 545-548.

Behrmann, M., & Plaut, D. C. (2013). Distributed circuits, not circumscribed centers, mediate visual recognition. *Trends in Cognitive Sciences, 17*(5), 210-219.

Bentin, S., Mouchetant-Rostaing, Y., Giard, M., Echallier, J., & Pernier, J. (1999). ERP manifestations of processing printed words at different psycholinguistic levels: time course and scalp distribution. *Journal of Cognitive Neuroscience, 11*(3), 235-260.

Binder, J. R., Medler, D. A., Westbury, C. F., Liebenthal, E., & Buchanan, L. (2006). Tuning of the human left fusiform gyrus to sublexical orthographic structure. *NeuroImage, 33*(2), 739-748. doi: 10.1016/j.neuroimage.2006.06.053

Boets, B., Op de Beeck, H. P., Vandermosten, M., Scott, S. K., Gillebert, C. R., Mantini, D., . . . Ghesquiere, P. (2013). Intact but less accessible phonetic representations in adults with dyslexia. *Science, 342*(6163), 1251-1254. doi: 10.1126/science.1244333

Borowsky, R., and Besner, D. (2006). Parallel distributed processing and lexical-semantic effects in visual word recognition: Are a few stages necessary? Psychological Review, 113, 181-193.

Bowers, J. S., Vigliocco, G., & Haan, R. (1998). Orthographic, phonological, and articulatory contributions to masked letter and word priming. *Journal of Experimental Psychology: Human Perception and Performance, 24*(6), 1705.

Brem, S., Bach, S., Kucian, K., Guttorm, T. K., Martin, E., Lyytinen, H., . . . Richardson, U. (2010). Brain sensitivity to print emerges when children learn letter-speech sound correspondences. *Proc Natl Acad Sci U S A, 107*(17), 7939-7944. doi: 10.1073/pnas.0904402107

Brunswick, N., McCrory, E., Price, C., Frith, C., & Frith, U. (1999). Explicit and implicit processing of words and pseudowords by adult developmental dyslexics. *Brain, 122*(10), 1901-1917.

Carreiras, M., Alvarez, C. J., & De Vega, M. (1993). Syllable frequency and visual word recognition in Spanish. *Journal of Memory and Language, 32*(6), 766.

Carreiras, M., Armstrong, B. C., Perea, M., & Frost, R. (2014). The what, when, where, and how of visual word recognition. *Trends in Cognitive Sciences, 18*(2), 90-98.

Carreiras, M., Dunabeitia, J. A., & Molinaro, N. (2009). Consonants and vowels contribute differently to visual word recognition: ERPs of relative position priming. *Cereb Cortex, 19*(11), 2659-2670. doi: 10.1093/cercor/bhp019

Carreiras, M., Dunabeitia, J. A., & Molinaro, N. (2012). Reading Words and Sentences in Spanish *The Handbook of Hispanic Linguistics* (J.I. Hualde, A. Olarrea & O'Rourke ed., pp. 803-826): Wiley-Blackwell.

Carreiras, M., Dunabeitia, J. A., & Perea, M. (2007). READING WORDS, NUMB3R5 and $YMB0L$. *Trends Cogn Sci, 11*(11), 454-455; author reply 456-457. doi: 10.1016/j.tics.2007.08.007

Carreiras, M., Ferrand, L., Grainger, J., & Perea, M. (2005). Sequential effects of phonological priming in visual word recognition. *Psychol Sci, 16*(8), 585-589. doi: 10.1111/j.1467-9280.2005.01579.x

Carreiras, M., Gillon-Dowens, M., Vergara, M., & Perea, M. (2009). Are vowels and consonants processed differently? Event-related potential evidence with a delayed letter paradigm. *J Cogn Neurosci, 21*(2), 275-288. doi: 10.1162/jocn.2008.21023

Carreiras, M., Mechelli, A., Estevez, A., & Price, C. J. (2007). Brain activation for lexical decision and reading aloud: two sides of the same coin? *J Cogn Neurosci, 19*(3), 433-444. doi: 10.1162/jocn.2007.19.3.433

Carreiras, M., Monahan, P. J., Lizarazu, M., Duñabeitia, J. A., & Molinaro, N. (2015). Numbers are not like words: Different pathways for literacy and numeracy. *NeuroImage, 118*, 79-89.

Carreiras, M., & Perea, M. (2002). Masked priming effects with syllabic neighbors in a lexical decision task. *J Exp Psychol Hum Percept Perform, 28*(5), 1228-1242.

Carreiras, M., Perea, M., Gil-Lopez, C., Abu Mallouh, R., & Salillas, E. (2013). Neural Correlates of Visual versus Abstract Letter Processing in Roman and Arabic Scripts. *J Cogn Neurosci*. doi: 10.1162/jocn_a_00438

Carreiras, M., Perea, M., Vergara, M., & Pollatsek, A. (2009). The time course of orthography and phonology: ERP correlates of masked priming effects in Spanish. *Psychophysiology, 46*(5), 1113-1122. doi: 10.1111/j.1469-8986.2009.00844.x

Carreiras, M., Quiñones, I., Hernández-Cabrera, J. A., & Duñabeitia, J. A. (2015). Orthographic coding: brain activation for letters, symbols, and digits. *Cerebral Cortex*, bhu163.

Carreiras, M., Quiñones, I., Mancini, S., Hernández-Cabrera, J. A., & Barber, H. (2015). Verbal and nominal agreement: An fMRI study. *NeuroImage, 120*, 88-103.

Carreiras, M., Seghier, M. L., Baquero, S., Estevez, A., Lozano, A., Devlin, J. T., & Price, C. J. (2009). An anatomical signature for literacy. *Nature, 461*(7266), 983-986. doi: 10.1038/nature08461

Carreiras, M., Vergara, M., & Barber, H. (2005). Early event-related potential effects of syllabic processing during visual word recognition. *J Cogn Neurosci, 17*(11), 1803-1817. doi: 10.1162/089892905774589217

Cohen, L., & Dehaene, S. (2000). Calculating without reading: Unsuspected residual abilities in pure alexia. *Cognitive neuropsychology, 17*(6), 563-583.

Cohen, L., & Dehaene, S. (2004). Specialization within the ventral stream: the case for the visual word form area. *NeuroImage, 22*(1), 466-476. doi: 10.1016/j.neuroimage.2003.12.049

Cohen, L., Dehaene, S., Chochon, F., Lehericy, S., & Naccache, L. (2000). Language and calculation within the parietal lobe: a combined cognitive, anatomical and fMRI study. *Neuropsychologia, 38*(10), 1426-1440.

Cohen, L., Dehaene, S., Naccache, L., Lehericy, S., Dehaene-Lambertz, G., Henaff, M. A., & Michel, F. (2000). The visual word form area: spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain, 123 (Pt 2)*, 291-307.

Cohen, L., Henry, C., Dehaene, S., Martinaud, O., Lehericy, S., Lemer, C., & Ferrieux, S. (2004). The pathophysiology of letter-by-letter reading. *Neuropsychologia, 42*(13), 1768-1780. doi: 10.1016/j.neuropsychologia.2004.04.018

Cohen, L., Lehericy, S., Chochon, F., Lemer, C., Rivaud, S., & Dehaene, S. (2002). Language-specific tuning of visual cortex? Functional properties of the Visual Word Form Area. *Brain, 125*(Pt 5), 1054-1069.

Cohen, L., Lehericy, S., Henry, C., Bourgeois, M., Larroque, C., Sainte-Rose, C., . . . Hertz-Pannier, L. (2004). Learning to read without a left occipital lobe: right-hemispheric shift of visual word form area. *Ann Neurol, 56*(6), 890-894. doi: 10.1002/ana.20326

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review, 108*(1), 204.

Conrad, M., Carreiras, M., Tamm, S., & Jacobs, A. M. (2009). Syllables and bigrams: orthographic redundancy and syllabic units affect visual word recognition at different processing levels. *J Exp Psychol Hum Percept Perform, 35*(2), 461-479. doi: 10.1037/a0013480

Cornelissen, P. L., Kringelbach, M. L., Ellis, A. W., Whitney, C., Holliday, I. E., & Hansen, P. C. (2009). Activation of the left inferior frontal gyrus in the first 200 ms of reading: evidence from magnetoencephalography (MEG). *PLoS One, 4*(4), e5359. doi: 10.1371/journal.pone.0005359

Christiansen, M. H., & Curtin, S. (1999). Transfer of learning: rule acquisition or statistical learning? *Trends in Cognitive Sciences, 3*(8), 289-290.

Christianson, K., Johnson, R. L., & Rayner, K. (2005). Letter transpositions within and across morphemes. *J Exp Psychol Learn Mem Cogn, 31*(6), 1327-1339. doi: 10.1037/0278-7393.31.6.1327

Christopher, M. E., Miyake, A., Keenan, J. M., Pennington, B., DeFries, J. C., Wadsworth, S. J., ... & Olson, R. K. (2012). Predicting word reading and comprehension with executive function and speed measures across development: a latent variable analysis. Journal of Experimental Psychology: General, 141(3), 470

Das, T., Padakannaya, P., Pugh, K. R., & Singh, N. C. (2011). Neuroimaging reveals dual routes to reading in simultaneous proficient readers of two orthographies. *NeuroImage*, *54*(2), 1476-1487. doi: 10.1016/j.neuroimage.2010.09.022

Davis, C. J. (2010). The spatial coding model of visual word identification. *Psychological Review, 117*(3), 713.

Dehaene, S., & Cohen, L. (2011). The unique role of the visual word form area in reading. *Trends Cogn Sci, 15*(6), 254-262. doi: 10.1016/j.tics.2011.04.003

Dehaene, S., Cohen, L., Sigman, M., & Vinckier, F. (2005). The neural code for written words: a proposal. *Trends in Cognitive Sciences, 9*(7), 335-341.

Dehaene, S., & Changeux, J. P. (2005). Ongoing spontaneous activity controls access to consciousness: a neuronal model for inattentional blindness. *PLoS Biol, 3*(5), e141. doi: 10.1371/journal.pbio.0030141

Diependaele, K., Duñabeitia, J. A., Morris, J., & Keuleers, E. (2011). Fast morphological effects in first and second language word recognition. *Journal of Memory and Language, 64*(4), 344-358.

Dimitropoulou, M., Duñabeitia, J. A., & Carreiras, M. (2011). Phonology by itself: Masked phonological priming effects with and without orthographic overlap. *J Cogn Psychol (Hove), 23*(2), 185-203.

Dunabeitia, J. A., & Carreiras, M. (2011). The relative position priming effect depends on whether letters are vowels or consonants. *J Exp Psychol Learn Mem Cogn, 37*(5), 1143-1163. doi: 10.1037/a0023577

Dunabeitia, J. A., Cholin, J., Corral, J., Perea, M., & Carreiras, M. (2010). SYLLABARIUM: an online application for deriving complete statistics for Basque and Spanish orthographic syllables. *Behav Res Methods, 42*(1), 118-125. doi: 10.3758/BRM.42.1.118

Dunabeitia, J. A., Molinaro, N., & Carreiras, M. (2011). Through the looking-glass: mirror reading. *NeuroImage, 54*(4), 3004-3009. doi: 10.1016/j.neuroimage.2010.10.079

Dunabeitia, J. A., Perea, M., & Carreiras, M. (2007). Do transposed-letter similarity effects occur at a morpheme level? Evidence for morpho-orthographic decomposition. *Cognition, 105*(3), 691-703. doi: 10.1016/j.cognition.2006.12.001

Dunabeitia, J. A., Perea, M., & Carreiras, M. (2009). There is no clam with coats in the calm coast: delimiting the transposed-letter priming effect. *Q J Exp Psychol (Hove), 62*(10), 1930-1947. doi: 10.1080/17470210802696070

Dunbar, K., & MacLeod, C. M. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance, 10*(5), 622.

Duñabeitia, J. A., Dimitropoulou, M., Estévez, A., & Carreiras, M. (2013). The Influence of Reading Expertise in Mirror-Letter Perception: Evidence From Beginning and Expert Readers. *Mind, Brain, and Education*, *7*(2), 124-135.

Duñabeitia, J. A., Molinaro, N., Laka, I., Estévez, A., & Carreiras, M. (2009). N250 effects for letter transpositions depend on lexicality: 'casual' or 'causal'? *NeuroReport, 20*(4), 381-387. doi: 10.1097/WNR.0b013e3283249b1c

Duñabeitia, J. A., Perea, M., & Carreiras, M. (2009). Eye movements when reading words with $YMβOL$ and NUM83R5: There is a cost. *Visual Cognition, 17*(5), 617-631. doi: 10.1080/13506280902764489

Duñabeitia, J. A., Perea, M., & Carreiras, M. (2014). Revisiting letter transpositions within and across morphemic boundaries. *Psychonomic Bulletin & Review, 21*(6), 1557-1575.

Eulitz, C., Eulitz, H., Maess, B., Cohen, R., Pantev, C., & Elbert, T. (2000). Magnetic brain activity evoked and induced by visually presented words and nonverbal stimuli. *Psychophysiology, 37*(04), 447-455.

Feldman, L. B., O'Connor, P. A., & del Prado Martín, F. M. (2009). Early morphological processing is morphosemantic and not simply morpho-orthographic: A violation of form-then-meaning accounts of word recognition. *Psychonomic Bulletin & Review, 16*(4), 684-691.

Ferrand, L., & Grainger, J. (1992). Phonology and orthography in visual word recognition: Evidence from masked non-word priming. *Quarterly Journal of Experimental Psychology: Section A, 45*(3), 353-372.

Ferrand, L., & Grainger, J. (1993). The time course of orthographic and phonological code activation in the early phases of visual word recognition. *Bulletin of the psychonomic society, 31*(2), 119-122.

Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for 'top-down'effects. *Behavioral and Brain Sciences*, 1-72.

Forster, K. I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *The Quarterly Journal of Experimental Psychology, 39*(2), 211-251.

Frank, S. L. (2006). Learn more by training less: systematicity in sentence processing by recurrent networks. *Connection Science, 18*(3), 287-302.

Frost, R. (2012). A universal approach to modeling visual word recognition and reading: Not only possible, but also inevitable. *Behavioral and Brain Sciences, 35*(05), 310-329.

Frost, R., Armstrong, B. C., Siegelman, N., & Christiansen, M. H. (2015). Domain generality versus modality specificity: the paradox of statistical learning. *Trends in Cognitive Sciences, 19*(3), 117-125.

Gerken, L. (2006). Decisions, decisions: Infant language learning when multiple generalizations are possible. *Cognition, 98*(3), B67-B74.

Giraud, A. L., & Ramus, F. (2013). Neurogenetics and auditory processing in developmental dyslexia. *Curr Opin Neurobiol, 23*(1), 37-42. doi: 10.1016/j.conb.2012.09.003

Gomez, P., Ratcliff, R., & Perea, M. (2008). The overlap model: a model of letter position coding. *Psychological Review, 115*(3), 577.

Goswami, U. (2011). A temporal sampling framework for developmental dyslexia. *Trends Cogn Sci, 15*(1), 3-10. doi: 10.1016/j.tics.2010.10.001

Goswami, U., & Ziegler, J. C. (2006). Fluency, phonology and morphology: a response to the commentaries on becoming literate in different languages. *Dev Sci, 9*(5), 451-453.

Grainger, J. (2008). Cracking the orthographic code: An introduction. *Language and Cognitive Processes, 23*(1), 1-35. doi: 10.1080/01690960701578013

Grainger, J., & Ferrand, L. (1996). Masked orthographic and phonological priming in visual word recognition and naming: Cross-task comparisons. *Journal of Memory and Language, 35*(5), 623-647.

Grainger, J., Granier, J. P., Farioli, F., Van Assche, E., & van Heuven, W. J. (2006). Letter position information and printed word perception: the relative-position priming constraint. *J Exp Psychol Hum Percept Perform, 32*(4), 865-884. doi: 10.1037/0096-1523.32.4.865

Grainger, J., & Holcomb, P. J. (2009a). An ERP investigation of orthographic priming with relative-position and absolute-position primes. *Brain Res, 1270,* 45-53. doi: 10.1016/j.brainres.2009.02.080

Grainger, J., & Holcomb, P. J. (2009b). Watching the Word Go by: On the Time-course of Component Processes in Visual Word Recognition. *Lang Linguist Compass, 3*(1), 128-156. doi: 10.1111/j.1749-818X.2008.00121.x

Grainger, J., & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: a multiple read-out model. *Psychological Review, 103*(3), 518.

Grainger, J., Kiyonaga, K., & Holcomb, P. J. (2006). The time course of orthographic and phonological code activation. *Psychol Sci, 17*(12), 1021-1026. doi: 10.1111/j.1467-9280.2006.01821.x

Grainger, J., Rey, A., & Dufau, S. (2008). Letter perception: from pixels to pandemonium. *Trends Cogn Sci, 12*(10), 381-387. doi: 10.1016/j.tics.2008.06.006

Grainger, J., & Ziegler, J. C. (2011). A dual-route approach to orthographic processing. *Front Psychol, 2,* 54. doi: 10.3389/fpsyg.2011.00054

Gregory, E., & McCloskey, M. (2010). Mirror-image confusions: Implications for representation and processing of object orientation. *Cognition, 116*(1), 110-129. doi: 10.1016/j.cognition.2010.04.005

Guerrera, C., & Forster, K. (2008). Masked form priming with extreme transposition. *Language and Cognitive Processes, 23*(1), 117-142. doi: 10.1080/01690960701579722

Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., & Lounasmaa, O. V. (1993). Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of modern Physics, 65*(2), 413.

Hoeft, F., Meyler, A., Hernandez, A., Juel, C., Taylor-Hill, H., Martindale, J. L., . . . Gabrieli, J. D. (2007). Functional and morphometric brain dissociation between dyslexia and reading ability. *Proc Natl Acad Sci U S A, 104*(10), 4234-4239. doi: 10.1073/pnas.0609399104

Holcomb, P. J., & Grainger, J. (2006). On the time course of visual word recognition: an event-related potential investigation using masked repetition priming. *J Cogn Neurosci, 18*(10), 1631-1643. doi: 10.1162/jocn.2006.18.10.1631

Hutzler, F., Bergmann, J., Conrad, M., Kronbichler, M., Stenneken, P., & Jacobs, A. M. (2004). Inhibitory effects of first syllable-frequency in lexical decision: an event-related potential study. *Neurosci Lett, 372*(3), 179-184. doi: 10.1016/j.neulet.2004.07.050

Ijalba, E., & Obler, L. K. (2015). First language grapheme-phoneme transparency effects in adult second-language learning. *Reading in a Foreign Language, 27*(1), 47.

Jacobs, A. M., & Grainger, J. (1991). Automatic letter priming in an alphabetic decision task. *Perception & Psychophysics, 49*(1), 43-52.

Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: A connectionist model. *Proceedings of the National Academy of Sciences, 96*(13), 7592-7597.

Jobard, G., Crivello, F., & Tzourio-Mazoyer, N. (2003). Evaluation of the dual route theory of reading: a metanalysis of 35 neuroimaging studies. *NeuroImage, 20*(2), 693-712. doi: 10.1016/s1053-8119(03)00343-4

Kherif, F., Josse, G., & Price, C. J. (2011). Automatic top-down processing explains common left occipito-temporal responses to visual words and objects. *Cereb Cortex, 21*(1), 103-114. doi: 10.1093/cercor/bhq063

Kim, W., Pitt, M. A., & Myung, J. I. (2013). How do PDP models learn quasiregularity? *Psychological Review, 120*(4), 903.

Kinoshita, S., & Kaplan, L. (2008). Priming of abstract letter identities in the letter match task. *The Quarterly Journal of Experimental Psychology, 61*(12), 1873-1885.

Kinoshita, S., Robidoux, S., Guilbert, D., & Norris, D. (2015). Context-dependent similarity effects in letter recognition. *Psychonomic Bulletin & Review, 22*(5), 1458-1464.

Klingberg, T., Hedehus, M., Temple, E., Salz, T., Gabrieli, J. D., Moseley, M. E., & Poldrack, R. A. (2000). Microstructure of temporo-parietal white matter as a basis for reading ability: evidence from diffusion tensor magnetic resonance imaging. *Neuron, 25*(2), 493-500.

Kronbichler, M., Hutzler, F., Wimmer, H., Mair, A., Staffen, W., & Ladurner, G. (2004). The visual word form area and the frequency with which words are encountered: evidence from a parametric fMRI study. *NeuroImage, 21*(3), 946-953. doi: 10.1016/j.neuroimage.2003.10.021

Laszlo, S., & Armstrong, B. C. (2014). PSPs and ERPs: Applying the dynamics of post-synaptic potentials to individual units in simulation of temporally extended Event-Related Potential reading data. *Brain and Language, 132*, 22-27.

Laszlo, S., & Federmeier, K. D. (2011). The N400 as a snapshot of interactive processing: Evidence from regression analyses of orthographic neighbor and lexical associate effects. *Psychophysiology, 48*(2), 176-186.

Laszlo, S., & Plaut, D. C. (2012). A neurally plausible parallel distributed processing model of event-related potential word reading data. *Brain Lang, 120*(3), 271-281. doi: 10.1016/j.bandl.2011.09.001

Lerner, I., Armstrong, B. C., & Frost, R. (2014). What can we learn from learning models about sensitivity to letter-order in visual word recognition? *Journal of Memory and Language, 77*, 40-58.

Lizarazu, M., Lallier, M., Molinaro, N., Bourguignon, M., Paz-Alonso, P. M., Lerma-Usabiaga, G., & Carreiras, M. (2015). Developmental evaluation of atypical auditory sampling in dyslexia: Functional and structural evidence. *Hum Brain Mapp, 36*(12), 4986-5002.

Lobier, M., Peyrin, C., Le Bas, J. F., & Valdois, S. (2012). Pre-orthographic character string processing and parietal cortex: a role for visual attention in reading? *Neuropsychologia, 50*(9), 2195-2204. doi: 10.1016/j.neuropsychologia.2012.05.023

Lukatela, G., & Turvey, M. (1990). Phonemic similarity effects and prelexical phonology. *Memory & Cognition, 18*(2), 128-152.

Lupker, S. J., Perea, M., & Davis, C. J. (2008). Transposed-letter effects: Consonants, vowels and letter frequency. *Language and Cognitive Processes, 23*(1), 93-116. doi: 10.1080/01690960701579714

Mareschal, D., & French, R. (in press) TRACX2: A connectionist autoencoder using graded chunks to model infant visual statistical learning. Philosophical Transactions B.

Massol, S., Carreiras, M., & Duñabeitia, J. A. (2016). Consonantal overlap effects in a perceptual matching task. *Experimental Brain Research*, 1-16.

Maurer, U., Brandeis, D., & McCandliss, B. D. (2005). Fast, visual specialization for reading in English revealed by the topography of the N170 ERP response. *Behav Brain Funct, 1*, 13. doi: 10.1186/1744-9081-1-13

Maurer, U., Brem, S., Bucher, K., & Brandeis, D. (2005). Emerging neurophysiological specialization for letter strings. *Journal of Cognitive Neuroscience, 17*(10), 1532-1552.

Maurer, U., Han, S., & McCandliss, B. (2005). *Fast visual specialization for print as recorded with high-density.* Paper presented at the Journal of Cognitive Neuroscience.

McClelland, J. L. (2009). The place of modeling in cognitive science. *Top Cogn Sci, 1*(1), 11-38.

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences, 14*(8), 348-356.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*(1), 1-86.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review, 88*(5), 375.

McClelland, J. L., St. John, M., & Taraban, R. (1989). Sentence comprehension: A parallel distributed processing approach. *Language and Cognitive Processes, 4*(3-4), SI287-SI335.

Molinaro, N., Duñabeitia, J. A., Marìn-Gutièrrez, A., & Carreiras, M. (2010). From numbers to letters: feedback regularization in visual word recognition. *Neuropsychologia, 48*(5), 1343-1355. doi: 10.1016/j.neuropsychologia.2009.12.037

Molinaro, N., Lizarazu, M., Lallier, M., Bourguignon, M., & Carreiras, M. (2016). Out-of-synchrony speech entrainment in developmental dyslexia. *Hum Brain Mapp*.

Mueller, S. T., & Weidemann, C. T. (2012). Alphabetic letter identification: Effects of perceivability, similarity, and bias. *Acta Psychol (Amst), 139*(1), 19-37.

Nation, K. (2009). Form–meaning links in the development of visual word recognition. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 364*(1536), 3665-3674.

Nickels, L., Biedermann, B., Coltheart, M., Saunders, S., & Tree, J. J. (2008). Computational modelling of phonological dyslexia: How does the DRC model fare? *Cognitive neuropsychology, 25*(2), 165-193.

Niogi, S. N., Mukherjee, P., & McCandliss, B. D. (2007). Diffusion tensor imaging segmentation of white matter structures using a Reproducible Objective Quantification Scheme (ROQS). *NeuroImage, 35*(1), 166-174.

O'Connor, R. E., & Forster, K. I. (1981). Criterion bias and search sequence bias in word recognition. *Memory & Cognition, 9*(1), 78-92.

Oliver, M., Carreiras, M., & Paz-Alonso, P. M. (2016). Functional Dynamics of Dorsal and Ventral Reading Networks in Bilinguals. *Cerebral Cortex*.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience, 8*(12), 976-987.

Paulesu, E., Frith, U., Snowling, M., Gallagher, A., Morton, J., Frackowiak, R. S., & Frith, C. D. (1996). Is developmental dyslexia a disconnection syndrome? *Brain, 119*(1), 143-157.

Paulesu, E., McCrory, E., Fazio, F., Menoncello, L., Brunswick, N., Cappa, S. F., . . . Lorusso, M. (2000). A cultural effect on brain function. *Nat Neurosci, 3*(1), 91-96.

Perea, M., & Carreiras, M. (1998). Effects of syllable frequency and syllable neighborhood frequency in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 24*(1), 134.

Perea, M., & Carreiras, M. (2006a). Do transposed-letter similarity effects occur at a prelexical phonological level? *Q J Exp Psychol (Hove), 59*(9), 1600-1613. doi: 10.1080/17470210500298880

Perea, M., & Carreiras, M. (2006b). Do Transposed-Letter Similarity Effects Occur at a Syllable Level? *Experimental Psychology (formerly "Zeitschrift für Experimentelle Psychologie"), 53*(4), 308-315. doi: 10.1027/1618-3169.53.4.308

Perea, M., Dunabeitia, J. A., & Carreiras, M. (2008a). R34d1ng W0rd5 W1th Numb3r5. *J Exp Psychol Hum Percept Perform, 34*(1), 237-241. doi: 10.1037/0096-1523.34.1.237

Perea, M., Dunabeitia, J. A., & Carreiras, M. (2008b). Transposed-letter priming effects for close versus distant transpositions. *Exp Psychol, 55*(6), 384-393.

Perea, M., Dunabeitia, J. A., Pollatsek, A., & Carreiras, M. (2009). Does the brain regularize digits and letters to the same extent? *Q J Exp Psychol (Hove), 62*(10), 1881-1888. doi: 10.1080/17470210902923374

Perea, M., & Estévez, A. (2008). Transposed-letter similarity effects in naming pseudowords: Evidence from children and adults. *European Journal of Cognitive Psychology, 20*(1), 33-46. doi: 10.1080/09541440701306941

Perea, M., & Fraga, I. (2006). Transposed-letter and laterality effects in lexical decision. *Brain Lang, 97*(1), 102-109. doi: 10.1016/j.bandl.2005.08.004

Perea, M., & Lupker, S. J. (2003). Does jugde activate COURT? Transposed-letter similarity effects in masked associative priming. *Memory & Cognition, 31*(6), 829-841.

Perea, M., & Lupker, S. J. (2004). Can CANISO activate CASINO? Transposed-letter similarity effects with nonadjacent letter positions. *Journal of Memory and Language, 51*(2), 231-246. doi: 10.1016/j.jml.2004.05.005

Perea, M., Mallouh, R. A., & Carreiras, M. (2010). The search for an input-coding scheme: transposed-letter priming in Arabic. *Psychon Bull Rev, 17*(3), 375-380. doi: 10.3758/PBR.17.3.375

Perea, M., Moret-Tatay, C., & Panadero, V. (2011). Suppression of mirror generalization for reversible letters: Evidence from masked priming. *Journal of Memory and Language, 65*(3), 237-246.

Perea, M., Palti, D., & Gomez, P. (2012). Associative priming effects with visible, transposed-letter nonwords: JUGDE facilitates COURT. *Attention, Perception, & Psychophysics, 74*(3), 481-488.

Perea, M., Vergara-Martínez, M., & Gomez, P. (2015). Resolving the locus of cAsE aLtErNaTiOn effects in visual word recognition: Evidence from masked priming. *Cognition, 142,* 39-43.

Peressotti, F., & Grainger, J. (1999). The role of letter identity and letter position in orthographic priming. *Perception & Psychophysics, 61*(4), 691-706.

Perry, C., Ziegler, J. C., & Zorzi, M. (2010). Beyond single syllables: large-scale modeling of reading aloud with the Connectionist Dual Process (CDP++) model. *Cogn Psychol, 61*(2), 106-151. doi: 10.1016/j.cogpsych.2010.04.001

Perry, C., Ziegler, J. C., & Zorzi, M. (2014). When silent letters say more than a thousand words: An implementation and evaluation of CDP++ in French. *Journal of Memory and Language, 72,* 98-115.

Petit, J. P., Midgley, K. J., Holcomb, P. J., & Grainger, J. (2006). On the time course of letter perception: a masked priming ERP investigation. *Psychon Bull Rev, 13*(4), 674-681.

Peyrin, C., Lallier, M., Demonet, J. F., Pernet, C., Baciu, M., Le Bas, J. F., & Valdois, S. (2012). Neural dissociation of phonological and visual attention span disorders in developmental dyslexia: FMRI evidence from two case reports. *Brain Lang, 120*(3), 381-394. doi: 10.1016/j.bandl.2011.12.015

Pinker, S., & Ullman, M. T. (2002). The past and future of the past tense. *Trends in Cognitive Sciences, 6*(11), 456-463.

Plaut, D. C. (1997). Structure and function in the lexical system: Insights from distributed models of word reading and lexical decision. *Language and Cognitive Processes*, *12*(5-6), 765-806.

Plaut, D. C. and Booth, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. Psychological Review, 107, 786-823.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: computational principles in quasi-regular domains. *Psychol Rev, 103*(1), 56-115.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: computational principles in quasi-regular domains. *Psychological Review, 103*(1), 56.

Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive neuropsychology, 10*(5), 377-500.

Pollatsek, A., Perea, M., & Carreiras, M. (2005). Does conal prime CANAL more than cinal? Masked phonological priming effects in Spanish with the lexical decision task. *Mem Cognit, 33*(3), 557-565.

Preston, J. L., Frost, S. J., Mencl, W. E., Fulbright, R. K., Landi, N., Grigorenko, E., . . . Pugh, K. R. (2010). Early and late talkers: school-age language, literacy and neurolinguistic differences. *Brain, 133*(Pt 8), 2185-2195. doi: 10.1093/brain/awq163

Price, C. J. (2012). A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. *NeuroImage, 62*(2), 816-847. doi: 10.1016/j.neuroimage.2012.04.062

Price, C. J., & Devlin, J. T. (2011). The interactive account of ventral occipitotemporal contributions to reading. *Trends Cogn Sci, 15*(6), 246-253. doi: 10.1016/j.tics.2011.04.001

Proctor, R. W. (1981). A unified theory for matching-task phenomena. *Psychological Review, 88*(4), 291.

Pugh, K. R., Frost, S. J., Sandak, R., Landi, N., Moore, D., Della Porta, G., . . . Mencl, W. E. (2010). Mapping the word reading circuitry in skilled and disabled readers. *The neural basis of reading*, 281-305.

Pugh, K. R., Mencl, W. E., Jenner, A. R., Katz, L., Frost, S. J., Lee, J. R., . . . Shaywitz, B. A. (2000). Functional neuroimaging studies of reading and reading disability (developmental dyslexia). *Ment Retard Dev Disabil Res Rev, 6*(3), 207-213. doi: 10.1002/1098-2779(2000)6:3<207::AID-MRDD8>3.0.CO;2-P

Ramus, F. (2003). Developmental dyslexia: specific phonological deficit or general sensorimotor dysfunction? *Curr Opin Neurobiol, 13*(2), 212-218.

Ramus, F., Pidgeon, E., & Frith, U. (2003). The relationship between motor control and phonology in dyslexic children. *J Child Psychol Psychiatry, 44*(5), 712-722.

Ramus, F., Rosen, S., Dakin, S. C., Day, B. L., Castellote, J. M., White, S., & Frith, U. (2003). Theories of developmental dyslexia: insights from a multiple case study of dyslexic adults. *Brain, 126*(Pt 4), 841-865.

Rastle, K., & Brysbaert, M. (2006). Masked phonological priming effects in English: Are they real? Do they matter? *Cognitive Psychology, 53*(2), 97-145.

Rastle, K., & Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes, 23*(7-8), 942-971. doi: 10.1080/01690960802069730

Rastle, K., Davis, M. H., & New, B. (2004). The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin & Review, 11*(6), 1090-1098.

Rayner, K., White, S. J., Johnson, R. L., & Liversedge, S. P. (2006). Raeding wrods with jubmled lettres there is a cost. *Psychological Science, 17*(3), 192-193.

Reicher, G. M. (1969). Perceptual recognition as a function of meaninfulness of stimulus material. *J Exp Psychol, 81*(2), 275-280.

Reilhac, C., Peyrin, C., Demonet, J. F., & Valdois, S. (2013). Role of the superior parietal lobules in letter-identity processing within strings: FMRI evidence from skilled and dyslexic readers. *Neuropsychologia, 51*(4), 601-612. doi: 10.1016/j.neuropsychologia.2012.12.010

Reingold, E. M., Yang, J., & Rayner, K. (2010). The time course of word frequency and case alternation effects on fixation times in reading: Evidence for lexical control of eye movements. *Journal of Experimental Psychology: Human Perception and Performance, 36*(6), 1677.

Richlan, F., Kronbichler, M., & Wimmer, H. (2011). Meta-analyzing brain dysfunctions in dyslexic children and adults. *NeuroImage, 56*(3), 1735-1742. doi: 10.1016/j.neuroimage.2011.02.040

Richlan, F., Kronbichler, M., & Wimmer, H. (2012). Structural abnormalities in the dyslexic brain: A meta-analysis of voxel-based morphometry studies. *Hum Brain Mapp*. doi: 10.1002/hbm.22127

Rohde, D. L. (2002). *A connectionist model of sentence comprehension and production.* Carnegie Mellon University Pittsburgh, PA.

Rosazza, C., Cai, Q., Minati, L., Paulignan, Y., & Nazir, T. A. (2009). Early involvement of dorsal and ventral pathways in visual word recognition: an ERP study. *Brain Res, 1272*, 32-44. doi: 10.1016/j.brainres.2009.03.033

Rueckl, J. G., Paz-Alonso, P. M., Molfese, P. J., Kuo, W.-J., Bick, A., Frost, S. J., . . . Duñabeitia, J. A. (2015). Universal brain signature of proficient reading: Evidence from four contrasting languages. *Proceedings of the National Academy of Sciences, 112*(50), 15510-15515.

Rueckl, J. G., & Rimzhim, A. (2011). On the interaction of letter transpositions and morphemic boundaries. *Language and Cognitive Processes, 26*(4-6), 482-508.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1988). Learning representations by back-propagating errors. *Cognitive modeling, 5*(3), 1.

Rumsey, J., Horwitz, B., Donohue, B., Nace, K., Maisog, J., & Andreason, P. (1997). Phonological and orthographic components of word recognition. A PET-rCBF study. *Brain, 120*(5), 739-759.

Rumsey, J., Horwitz, B., Donohue, B., Nace, K., Maisog, J., & Andreason, P. (1999). A functional lesion in developmental dyslexia: left angular gyral blood flow predicts severity. *Brain and Language, 70*(2), 187-204.

Rumsey, J. M., Andreason, P., Zametkin, A. J., Aquino, T., King, A. C., Hamburger, S. D., . . . Cohen, R. M. (1992). Failure to activate the left temporoparietal cortex in dyslexia: An oxygen 15 positron emission tomographic study. *Arch Neurol, 49*(5), 527-534.

Rumsey, J. M., Nace, K., Donohue, B., Wise, D., Maisog, J. M., & Andreason, P. (1997). A positron emission tomographic study of impaired word recognition and phonological processing in dyslexic men. *Arch Neurol, 54*(5), 562-573.

Schoonbaert, S., & Grainger, J. (2004). Letter position coding in printed word perception: Effects of repeated and transposed letters. *Language and Cognitive Processes, 19*(3), 333-367. doi: 10.1080/01690960344000198

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychol Rev, 96*(4), 523-568.

Seidenberg, M. S., & Plaut, D. C. (2014). Quasiregularity and its discontents: the legacy of the past tense debate. *Cognitive Science, 38*(6), 1190-1228.

Share, D. L. (2008). On the Anglocentricities of current reading research and practice: the perils of overreliance on an" outlier" orthography. *Psychological Bulletin, 134*(4), 584.

Shaywitz, S. E., Shaywitz, B. A., Pugh, K. R., Fulbright, R. K., Constable, R. T., Mencl, W. E., . . . Gore, J. C. (1998). Functional disruption in the organization of the brain for reading in dyslexia. *Proc Natl Acad Sci U S A, 95*(5), 2636-2641.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychol (Amst), 30*, 276-315.

Taft, M., & Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of verbal learning and verbal behavior, 14*(6), 638-647.

Tarkiainen, A., Cornelissen, P. L., & Salmelin, R. (2002). Dynamics of visual feature analysis and object-level processing in face versus letter-string perception. *Brain, 125*(Pt 5), 1125-1136.

Tarkiainen, A., Helenius, P., Hansen, P. C., Cornelissen, P. L., & Salmelin, R. (1999). Dynamics of letter string perception in the human occipitotemporal cortex. *Brain, 122 (Pt 11)*, 2119-2132.

Thesen, T., McDonald, C. R., Carlson, C., Doyle, W., Cash, S., Sherfey, J., . . . Halgren, E. (2012). Sequential then interactive processing of letters and words in the left fusiform gyrus. *Nat Commun, 3*, 1284. doi: 10.1038/ncomms2220

Thiebaut de Schotten, M., Cohen, L., Amemiya, E., Braga, L. W., & Dehaene, S. (2012). Learning to Read Improves the Structure of the Arcuate Fasciculus. *Cereb Cortex*. doi: 10.1093/cercor/bhs383

Tranel, D., Logan, C. G., Frank, R. J., & Damasio, A. R. (1997). Explaining category-related effects in the retrieval of conceptual and lexical knowledge for concrete entities: Operationalization and analysis of factors. *Neuropsychologia*, *35*(10), 1329-1339.

Velan, H., & Frost, R. (2007). Cambridge University versus Hebrew University: The impact of letter transposition on reading English and Hebrew. *Psychonomic Bulletin & Review, 14*(5), 913-918.

Velan, H., & Frost, R. (2009). Letter-transposition effects are not universal: The impact of transposing letters in Hebrew. *Journal of Memory and Language, 61*(3), 285-302.

Velan, H., & Frost, R. (2011). Words with and without internal structure: What determines the nature of orthographic and morphological processing? *Cognition, 118*(2), 141-156.

Vergara-Martinez, M., Perea, M., Marin, A., & Carreiras, M. (2011). The processing of consonants and vowels during letter identity and letter position assignment in visual-word recognition: an ERP study. *Brain Lang, 118*(3), 105-117. doi: 10.1016/j.bandl.2010.09.006

Wheat, K. L., Cornelissen, P. L., Frost, S. J., & Hansen, P. C. (2010). During visual word recognition, phonology is accessed within 100 ms and may be mediated by a speech production code: evidence from magnetoencephalography. *J Neurosci, 30*(15), 5229-5233. doi: 10.1523/JNEUROSCI.4448-09.2010

White, S. J., Johnson, R. L., Liversedge, S. P., & Rayner, K. (2008). Eye movements when reading transposed text: the importance of word-beginning letters. *Journal of Experimental Psychology: Human Perception and Performance, 34*(5), 1261.

Whitney, C. (2001). How the brain encodes the order of letters in a printed word: The SERIOL model and selective literature review. *Psychonomic Bulletin & Review, 8*(2), 221-243.

Whitney, C., & Berndt, R. S. (1999). A new model of letter string encoding: Simulating right neglect dyslexia. *Prog Brain Res, 121*, 143-163.

Woodhead, Z. V., Barnes, G. R., Penny, W., Moran, R., Teki, S., Price, C. J., & Leff, A. P. (2012). Reading Front to Back: MEG Evidence for Early Feedback Effects During Word Recognition. *Cereb Cortex*. doi: 10.1093/cercor/bhs365

Yang, J., McCandliss, B. D., Shu, H., & Zevin, J. D. (2009). Simulating language-specific and language-general effects in a statistical learning model of Chinese reading. Journal of Memory and Language, 61(2), 238-257.

Ziegler, J. C., Ferrand, L., Jacobs, A. M., Rey, A., & Grainger, J. (2000). Visual and phonological codes in letter and word recognition: Evidence from incremental priming. *The Quarterly Journal of Experimental Psychology: Section A, 53*(3), 671-692.

Ziegler, J. C., & Goswami, U. (2006). Becoming literate in different languages: Similar problems, different solutions. *Dev Sci, 9*(5), 429-436.

Figure Captions

Figure 1: Examples of different orthographic inputs.

Figure 2. The interactive activation model of visual word recognition. The model maps from early visual representations of line segments, through to individual letters, and finally, to words. Connections between layers (levels of representation) carry excitation, whereas connections within layers carry inhibition. Adapted from McClelland & Rumelhart (1981).

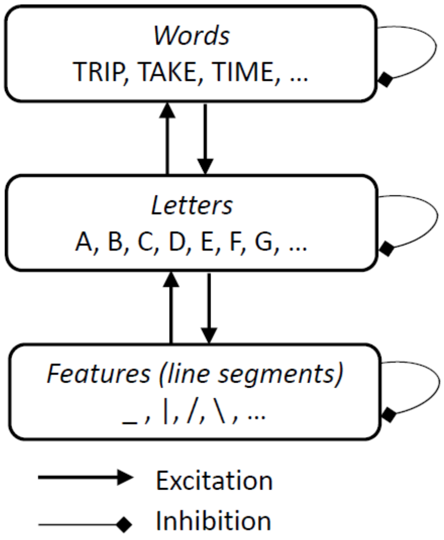

Figure 3. Depiction of the warping required to represent an exception word (*pint*) along with neighbouring words with regular pronunciations. There is some spill-over to nearby words (e.g., *tint*, *hint*) and to neighbouring nonwords (e.g., *kint*, *gint*). Explicit training on words leaves neighbouring regular word pronunciations intact. Figure adapted from Armstrong et al. (Accepted).

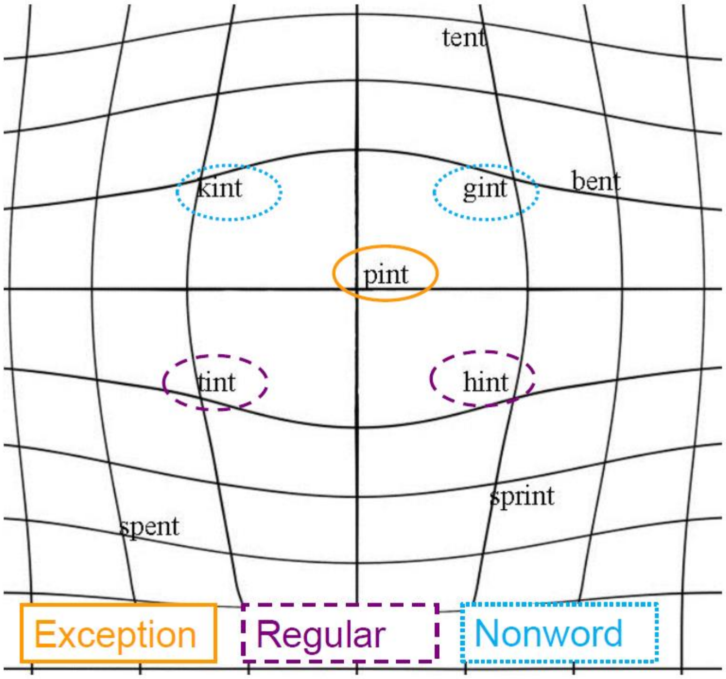

Figure 4. Surface rendering of the brain depicting neural activation in a lexical decision task. The activity maps plot activation for words and pseudowords compared to deciding ‘‘no’’ to false fonts and of reading aloud words and pseudowords compared to responding ‘‘false’’ to false fonts. All contrasts depicted at p < .05, corrected (from Carrreiras et al., 2007. JOCN).