



HHS Public Access

Author manuscript

Sci Stud Read. Author manuscript; available in PMC 2016 March 18.

Published in final edited form as:

Sci Stud Read. 2016 January 1; 20(1): 86–97. doi:10.1080/10888438.2015.1103741.

Towards a Theory of Variation in the Organization of the Word Reading System

Jay G. Rueckl

University of Connecticut and Haskins Laboratories

Abstract

The strategy underlying most computational models of word reading is to specify the organization of the reading system—its architecture and the processes and representations it employs—and to demonstrate that this organization would give rise to the behavior observed in word reading tasks. This approach fails to adequately address the variation in reading behavior observed across and within linguistic communities. Only computational models that incorporate learning can fully account for variation in organization. However, even extant learning models (e.g., the triangle model) must be extended if they are to fully account for variation in organization. The challenges associated with extending theories in this way are discussed.

One hallmark of the theoretical literature on skilled word reading is its emphasis on mechanism. Although some theoretical treatments have attempted to characterize word recognition in terms of broad theoretical principles (e.g. Frost, 1998) or in abstract rational terms (Norris, 2006), the more typical approach has been to characterize the *organization* of the reading system: the processes by which words are recognized, the representations over which these processes operate, and the computational architecture in which these processes occur. The most explicitly detailed theories of this sort are embodied in computational simulations (e.g., Coltheart et al., 2001; Grainger & Jacobs, 1996; McClelland & Rumelhart, 1981). Computational modeling forces theorists to make explicit aspects of their theory that might otherwise remain implicit and has allowed modelers to investigate subtle aspects of the word recognition process that might have otherwise been ignored. Simulation methods are particularly valuable when the theory is of sufficient complexity to render intuitive judgments of the model's behavior untrustworthy. Moreover, computational modeling has directly stimulated a substantial body of empirical research (e.g., Andrews & Scarrett, 1996; Pritchard et al., 2012; Seidenberg et al. 1994; Treiman et al., 2003). Thus, the computational modeling approach has been an important development in the study of word reading.

That being said, many computational models are rooted in the same scientific strategy as the “box-and-arrow” models that preceded them (e.g., Morton, 1969). Elsewhere I have referred to this strategy as “reverse engineering” (Rueckl, 2012; also see, for example, Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010). The application of this strategy begins with the identification of a circumscribed set of target phenomena. (In the case of word

recognition, these might include the effects of word frequency, letter transposition, and semantic priming on tasks such as lexical decision and naming.) A mechanistic account of these phenomena is derived; if it is a computational modeling account, the *organization* of the system (the hypothesized representations and processes) is described in sufficient detail so that the simulations of the model can be performed. The adequacy of the model is then demonstrated by showing that the hypothesized mechanism would generate the target phenomenon. (In box-and-arrow models, this demonstration is in the form of a verbal explication; for computational models, it takes the form a comparison of simulation results and empirical data.)

For the most part, theories of this sort have been geared towards providing a model that is representative of a typical reader. To be sure, such models have occasionally been used to capture individual differences (e.g., Ziegler et al., 2008). In general, however, in the reverse-engineering approach differences among skilled adult readers are typically ignored: Experimental data are usually characterized by a measure of central tendency, and variability around this measure is usually treated as observational noise. Moreover, because such theories provide a static snapshot of a reader's neurocognitive organization, they fail to provide a mechanistic account of the processes that give rise to differences among readers—a question of particular relevance for those, for example, who seek to understand, diagnose, and treat reading disability.

The purpose of this article is to consider the implications of taking the explanation of variability in organization as the primary goal of a computational theory. In the next section, cross-language differences are discussed to motivate this goal and to highlight the central role of plasticity in such a theory. Following this is an explication of one particular computational word reading model that incorporates learning: the triangle model (Seidenberg & McClelland, 1989; Plaut, McClelland, Seidenberg, & Patterson, 1996; Harm & Seidenberg, 2004). This model provides an illustration of how a computational model can provide a mechanistic account of both the processes by which a word is read and the processes by which these processes change over the course of learning—a requisite property for any theory meant to explain variation in reading organization. Notably, however, the triangle model is similar to models rooted in the reverse-engineering approach in that it has largely been used to model the 'typical' reader of a given population. (Important exceptions to this general claim are noted below). It is argued below that extending the triangle model to provide a general account of variation in reading organization requires the development of a broader theoretical framework. This broader framework is somewhat generic, and could likely accommodate alternatives to the triangle model (provided they incorporate a learning mechanism).

Organization as Explanadum

Theories rooted in the reverse-engineering approach take the organization of the reading system as a theoretical primitive. Questions about the processes by which that organization comes about are beyond the scope of such theories. For an illustration of the limitations of this approach, consider what cross-language research has revealed about the reading process.

Much of this research has been driven by the observation that writing systems differ in the manner in which the phonological structure of the words of a language is represented by their written forms. Some writing systems (e.g., Spanish, Serbo-Croatian) are highly transparent in that there is a nearly one-to-one mapping between these two domains. (That is, each letter represents a particular phoneme, and each phoneme is written by a particular letter.) In contrast to these *orthographically shallow* writing systems, in *orthographically deep systems* the orthographic-phonological mapping is much more ambiguous. (Examples include Chinese, which has extensive homophony, and unpointed Hebrew, in which the written forms are missing most of the vowels.) English represents an intermediate case: a high degree of regularity, coupled with irregularities such as the pronunciation of the vowel in PINT and the final consonant in COMB.

An extensive body of research has demonstrated that reading behavior differs systematically as a consequence of the orthographic depth of the language being read (e.g. Frost et al., 1987; Paulesu et al. 2010; Ziegler et al., 2010). This patterning of reading behavior has been attributed to differences in the relative contribution of two 'pathways' used in the reading process. On one pathway, a word's lexical (or semantic) representation is directly retrieved based on information about the orthographic structure of the input; on the other pathway, access to this lexical/semantic representation is mediated by a phonological representation that is assembled based on sub-lexical spelling-to-sound knowledge. According to the orthographic depth hypothesis (Frost et al., 1987), readers of shallow orthographies tend to rely more on the phonological pathway. Put another way, for readers of shallow orthographies, the *division of labor* (Harm & Seidenberg, 2004) between the phonological and lexical/semantic pathways is tilted more towards phonology.

This account is appealing in that it grounds an account of cross-language differences in a dual-pathway framework that has been used to explain a variety of other reading phenomena, including the effects of lexicality, frequency, and orthographic-phonological regularity (Coltheart et al., 2001), the consequence of brain damage on reading (Coltheart, 2006), and differences between typically developing and reading-disabled children (Ziegler et al., 2008), and indeed, computationally implemented variants of this framework have been used to model reading in several languages, including English (Coltheart et al., 2001) and German (Ziegler, Perry, & Coltheart, 2001).

On the other hand, the orthographic depth hypothesis is clearly an incomplete explanation. For one thing, the orthographic depth hypothesis is relatively vague about language-related differences in the organization of the reading system. It holds that properties of the writing system determine the relative strength of the phonological and lexical-semantic pathways, but like other box-and-arrow (verbal) accounts, it does not characterize the computational mechanisms determining the speed of these processes in any detail.

More importantly, as is typical of theories stemming from the reverse-engineering tradition, the orthographic depth hypothesis takes organization as a theoretical primitive and thus can do no more than stipulate how the organization of the reading system differs across populations. The orthographic depth hypothesis is almost surely correct in asserting that the nature of the writing system determines the relative importance of the phonological and

lexical-semantic pathways, but it is silent about questions of both theoretical and practical importance. For example, what is the process that aligns the organization of a child's reading system to the properties of the written language that child is exposed to? More broadly, why are some instructional methods more likely than others to engender an appropriate organization? Similarly, why do some children at risk for reading disability respond to a particular intervention while others do not?

To answer these questions, a theory must do more than explain how the organization of the reading system determines behavior; it must also explain *how* that organization comes about. Thus, organization has a dual role: It provides the explanation for certain phenomena, but it is also a phenomenon to be explained. In other words, organization is both explanation and explanandum.

Plasticity and Computational Models

As cross-language differences make clear, the organization of the reading system is at least in part experience-dependent. This implies that an explanation of variation in the organization of the reading system must incorporate a learning mechanism. Generally, because models of visual word recognition have typically been rooted in the reverse-engineering approach, they have rarely confronted the question of learning. However, since the introduction of the *triangle model* (Seidenberg & McClelland, 1989), learning has begun to figure more prominently in models of reading, both in the further development of the triangle model (e.g., Harm & Seidenberg, 1999, 2004; Plaut et al. 1996) and in models grounded in other perspectives (e.g. the CDP++ model, Perry, Ziegler, Zorzi, 2007, 2010; Ziegler, Perry, & Zorzi, 2014).¹

The triangle model is a theory of reading based on principles of the connectionist or PDP framework (Rumelhart et al., 1986, also see Elman et al., 1996). In this framework, cognitive systems are network composed of many simple, neuron-like processing units (*nodes*) that communicate by sending excitatory and inhibitory signals to one another. Each signal is weighted by the strength of the connection that it is sent across, and the state of each node (its *activation*) is a nonlinear function of the sum of these weighted signals. Like neural synapses, the connections in a network are plastic, and a learning algorithm is used to adjust their strengths (or *weights*) based on the interaction of the network and its task environment.

The triangle model assumes that the word reading system is a connectionist network specifically tasked with learning to read words. Seeing a word causes a flow of activation in this network; over time the network settles into a stable pattern of activation that serves as the representation of that word. The network is organized into distinct 'layers' (sets of nodes) responsible for representing the various linguistic properties (orthographic,

¹It is convenient to treat learning models as categorically distinct from the static models developed within the reverse-engineering approach, but the contrast is subtler than this. *All* computational models stipulate some theoretical primitives; in the triangle model these usually include the architecture, the activation and learning functions, and the input (orthographic) and output (phonological and semantic) representations. However, models often differ in the number and complexity of these primitives, the degree to which the primitives are domain specific, and the considerations used to justify the selection of the primitives. One key difference between learning and non-learning reading models is that in the latter domain-specific knowledge is taken as a primitive.

phonological, semantic) of the input. The connections among these layers are organized such that the triangle model is an instantiation of the dual-pathway framework, with distinct (but interacting) subnetworks mapping orthography to phonology and semantics. Computational implementations of the model have simulated a variety of aspects of reading. (See Plaut, 1999, Rueckl & Seidenberg, 2009, and Seidenberg, 2005, for reviews). One especially pertinent aspect of this work has been an increasingly deep understanding of how learning shapes the behavior of the reading network.

In the triangle model, each time a word is presented, the resulting pattern of activation is compared to a target pattern (a pattern representing the correct pronunciation and meaning of that word). A simple algorithm (*backpropagation*) determines how each connection weight will be changed given the difference between these patterns (Rumelhart et al. 1986). Because the changes to the weights are small (and must be, to avoid *catastrophic interference*, McClelland et al, 1995) the impact of any single learning event is also quite small. Nonetheless, because learning is incremental, over the course of many learning events the network becomes an increasingly skilled reader, and (all else being equal) the more frequently a word is encountered the better the network's performance on that word.

An important aspect of this learning mechanism is that learning generalizes. That is, even though the weight changes made on any particular learning event are specified so as to improve performance on subsequent encounters with that input, the fact that the network employs distributed representations means that these changes affect the response to other inputs as well. Whether the response to one word is affected by learning about another word depends on input similarity—how similarly the words are spelled. (The more similar, the greater the impact.) Whether performance on the transfer item is improved or harmed (whether transfer is seen as 'generalization' or 'interference') depends on output similarity—whether the words are alike phonologically or semantically.

In the triangle model, this tendency for generalization results in the network becoming sensitive to statistical regularities in the mappings between orthography, phonology, and semantics. An extensive body of research (e.g., Seidenberg & McClelland, 1989; Plaut & Gonnerman, 2000) has illuminated how the behavior of a network (and, by theory, of a reader) is shaped by these regularities. Initial simulations focused on the phonological and lexical/semantic pathways separately, revealing how a network becomes attuned to the regularities of the mapping performed by that pathway. Subsequently, more of the focus has been on the cooperative interaction of these pathways in word reading (Harm & Seidenberg, 2004; Welbourne, Woollams, Crisp, & Lambon Ralph, 2011). These investigations have revealed how the statistical properties of the orthographic-phonological and orthographic-semantic mappings shape the division of labor between the pathways that perform these mappings.

A key fact about the structure of the orthographic-phonological and orthographic-semantic mappings in most writing systems is that the orthographic-phonological mapping is much more systematic. For example, in English, most words that look alike sound alike. Thus, the word body -ILL is consistently mapped to /il/ (e.g., PILL, MILL, HILL). Similarly, the word body -INT is usually pronounced /int/ (as in MINT), although there is an exception to this

regularity (PINT). In contrast, there are far fewer regularities in the mapping from orthography to semantics—words that look alike are generally unrelated in meaning (e.g., BAKE, TAKE, and LAKE). Morphological structure does impose a certain amount of regularity on this mapping. (E.g. BAKE, BAKER, and BAKERY are similar in form and meaning.) However, morphological structure creates islands of regularity in a sea of arbitrary correspondences, whereas in the mapping from orthography to phonology, irregular pronunciations form small islands in a sea of systematicity.²

Harm & Seidenberg (2004) developed a number of methods to investigate how these properties of the orthographic-phonological and orthographic-semantic mappings shape the division of labor between their corresponding pathways. For example, they examined the behavior of the network after removing the weights connecting the orthographic units to either the semantic or the phonological units after different amounts of training. Early in training, the division of labor was strongly tilted towards phonology: removing the orthographic-phonological weights greatly impaired the networks performance, whereas removing the orthographic-semantic connections was far less disruptive. This pattern is a direct consequence of the difference in systematicity between the orthographic-phonological and orthographic-semantic mappings. Systematic mappings are much easier for a network to learn. Hence, the systematic orthographic-phonological mapping was mastered first, and the phonological pathway dominated the behavior of the network. With additional experience, the network continued to acquire knowledge of the orthographic-semantic mapping, and by the end of training could either be read by either pathway or could only be read when both pathways were intact. Thus, after sufficient training there was a more 'equitable' division of labor between the pathways.

These simulations suggest that the division of labor between phonology and semantics is a function of the disparity in the systematicity of the orthographic-phonological and orthographic-semantic mappings. This is, of course, consistent with the pattern of cross-language differences cited above that gave rise to the orthographic depth hypothesis. In line with this account, Yang, Shu, McCandliss, and Zevin (2012) trained a variant of the triangle model on either Chinese or English. As expected, the division of labor within these networks differed with a stronger tilt towards semantics in the network trained on Chinese, a difference that Yang et al. attributed to differences in the statistical properties of the writing system. More recently, Lerner, Armstrong, & Frost, (2014) made a similar point by linking the statistical properties of writing systems to another cross-language difference—the impact of letter transpositions.

Towards a Theory of Variation in the Organization of the Word Reading System

The discussion thus far has focused on two key points. The first is that organization plays a dual role in theory: It provides an explanation for some aspects of behavior but is also, itself,

²Just as the systematicity of the orthographic-phonological mapping varies across writing systems (as discussed in the previous section), so too does the systematicity of the orthographic-semantic mappings. See Plaut and Gonnerman (2000) for simulations exploring how this factor (termed *morphological richness*) influences a reading network's behavior.

a phenomenon that must be explained. The second is that plasticity is critical for explaining how a cognitive organization comes about and, in particular, how experience shapes that organization. It is noteworthy that because experience shapes organization, and because individuals differ in both their experiences and in 'constitutional' factors (e.g., system parameters that are set independently of reading experience), variation in organization is virtually inevitable. In the case of word reading, cross-language differences in organization illustrate this point. So, too, do changes in organization over the course of acquisition (e.g., Backman et al., 1984), differences associated with developmental dyslexia (e.g., Manis et al., 1996), and differences among individuals with a linguistic community (e.g., Yap et al., 2012).

As discussed in the previous section, the triangle model provides an example of a computational theory in which plasticity gives rise to variation in organization. It is important to note, however, that the triangle model is not the only extant computational model to incorporate learning (Perry, Ziegler, Zorzi, 2007, 2010; Ziegler, Perry, & Zorzi, 2014) and it seems likely that the number of such models will grow in the future³. It is also important to note that, like most other word reading models, the primary application of the triangle model has been (to a large extent) as a model of a typical reader (and for the most, a typical reader of English). That said, the triangle model has been used to address cross-language differences (Yang et al, 2012) and variation in organization associated with both developmental (Harm & Seidenberg, 1999) and acquired dyslexia (Dilkina, McClelland, & Plaut, 2008; Wellbourne et al., 2012) as well as environmental differences (Harm et al., 2003; Zevin & Seidenberg, 2006). Thus, although the triangle model is not a theory that takes as its primary goal the explanation of variation in the organization of the word reading system, these extensions provide some indication of what a theory of variation in organization might look like.

First, the theory should specify a set of control parameters that shape the organization (and hence the behavior) of the system.⁴ These parameters characterize the architecture of the network or act as constraints on the system's computational primitives. Different settings on the control parameters can give rise to different patterns of organization. For example, in an application of the triangle model to developmental dyslexia, Harm and Seidenberg (1999) demonstrated that both the kinds of reading deficits exhibited by a network as well as the severity of those deficits varied with the values of control parameters such as the amount of noise in the phonological system, the size of the reading network, and the magnitude of the 'weight changes' made on each learning event. Similarly, Welbourne et al. (2011) demonstrated that variation in the severity of damage to the phonological and semantic systems can give rise to variation in the pattern of deficits (and recovery) observed in acquired dyslexia.

³It is worth noting that in both the triangle model and the CDP++ models learning acts to optimize the associations among stipulated representations. However, in the triangle model learning plays the additional role of creating new internal ('hidden') representations that mediate these associations. This is of theoretical importance because one long-term goal of from the triangle model approach is to eliminate stipulated representations and to explain all representations as the product of the learning (Rueckl & Seidenberg, 200; Seidenberg, 2011.).

⁴The term 'control parameter' is likely used in many branches of science. In the present context, this use of this term is most closely related to the field of nonlinear dynamics. For discussions of the linkages between nonlinear dynamics and reading, see Rueckl, 2002, and Van Orden et al., 1999).

In the present context, much of the research on individual differences in reading can be understood as the investigation of the impact of control parameters on reading organization. Candidate parameters include endogenous factors specified at the level of genes (Landi et al., 2013), the brain (Pugh et al., 2014), or cognition (e.g., perceptual skill, Holyk & Pexman, 2004; Plaut & Booth, 2000), as well as exogenous factors such as the properties of a writing system (Frost, 2012) or the structure of the instructional curriculum (Harm et al., 2003). Importantly, an important goal for this research is not just to determine whether and how various parameters constrain organization, but also how parameters specified at different levels relate to one another. For example, one interesting hypothesis is that the noise parameter manipulated in Harm and Seidenberg's (2004) computational simulations reflects variation in the neurochemistry of certain parts of the brain, which may in turn reflect genetic constraints (Pugh, 2014),

It is also important to note that although there is a systematic relationship between a system's control variables and its behavior, this relationship is not necessarily fully deterministic. In part (and relatively uninterestingly) this can arise because either the system or the observations of that system are noisy. More importantly, for systems that learn, the behavior of the system can change even if the control parameters remain fixed, as demonstrated by developmental changes in the division of labor revealed in the Harm and Seidenberg (2004) simulations discussed above.⁵ Moreover, stochastic sampling of the environment can itself give rise to different patterns of behavioral change (Zevin & Seidenberg, 2006). Thus, it is important to distinguish between a system's control parameters and its organization. A system's organization is the proximate cause of its behavior. The organization is constrained by the control parameters, but it is not identical to them.⁶

A key question, then, is how to characterize the organization of the reading system. Here, the triangle model exemplifies a potentially useful construct. As with other connectionist networks, learning in the triangle model can be characterized as a search through its weight space, a geometric space such that each possible pattern of connectivity corresponds to a point in this space, with similar patterns of connectivity corresponding to nearby points in the space.⁷ The small weight changes that are made during a given learning event correspond to small moves within the network's weight space, and the accumulation of changes over the course of learning corresponds to a trajectory through its weight space.

At the level of individual networks (or, by theory, readers), it may be possible to equate a network's position in its weight space and its organization as a reading system. However, at the level of a population of individuals, this equivalence breaks down.⁸ What is needed is a way to describe possible organizations of the reading system in a way that is more transparently related to the behavioral consequences of that organization—a description of

⁵This is not to say that the control parameters *must* remain fixed. They may well change over time, and with interesting consequences. The point here is simply that for systems that learn, even if the control parameters remain constant, behavior (and the organization underlying it) does not.

⁶In the language of dynamical systems, the distinction being drawn here is the distinction between *control* and *state* (or *order*) parameters (Rueckl, 2002).

⁷Technically, the weight space is a high-dimensional space where each dimension corresponds to one connection and the network's position along this dimension corresponds to the weight of that connection.

the system at the task level. One possibility is to hypothesize that the reading system lives in an organizational space, such that each dimension corresponds to a dimension of variation among readers, nearby points in the space correspond to similar organizations, and trajectories through this space correspond to the changes in organization that occur as individuals learn to read.

The general structure of this framework is presented in Figure 1. For the purpose of illustration, the organizational space (in the center of the figure) is depicted as two-dimensional. One dimension corresponds to the division of labor between the phonological and lexical-semantic pathways, which has been hypothesized to vary both across (see above) and within (Strain & Herdman, 1999; Welbourne et al., 2011) linguistic communities. The second dimension corresponds to the reliance on orthographic units of differing grain sizes (e.g., single letters vs letter clusters such as *-int*). Like the division of labor, sensitivity to varying grain sizes has also been hypothesized to vary within (Treiman et al., 2006) and between (Ziegler & Goswami, 2005) linguistic communities.

Hypothetical developmental trajectories through this organizational space are depicted by the lines labeled A, B, and C. It has been hypothesized that the acquisition of reading skill in English is characterized by a shift in the division of labor towards the lexical/semantic pathway (Harm & Seidenberg, 2004; rightwards in the figure) and towards the use of larger orthographic units (Treiman et al., 2003; upwards in the figure). Thus, line A depicts what might be taken as the typical developmental trajectory for English; lines B & C depict other trajectories.

What gives rise to the differences among these trajectories? As argued above, a primary source of variation in the organization of reading is variation in the setting of the reading system's control parameters. Again for the purposes of illustration, Figure depicts two such parameters. Let us suppose that the endogenous parameter corresponds to the relative efficiency of two domain-general learning systems (e.g., the hippocampal and non-hippocampal systems; McClelland, McNaughton, & O'Reilly, 1995; Davis & Gaskell, 2009), and that these learning systems are differentially involved in the acquisition of different kinds of linguistic knowledge (with the hippocampal system particularly important for learning word meaning; Ullman, 2004). Moreover, let us also suppose that a greater contribution of semantics to reading (as might be the case for someone who is atypically reliant on hippocampal learning) results in less pressure for the reading system to discover large-grain regularities in the mapping from orthography to phonology.⁹ Lines A and B in Figure 1 illustrate these relationships. Relative to individual A, individual B is more hippocampal-dominant (rightward in the control space). As a consequence (given the above

⁸The reasons for this breakdown in equivalence are rather technical. First, the dimensionality of a network's weight space depends precisely on the number of processing units (nodes), and there is no reason to expect that this number is constant across individuals. Second, the relationship between positions in weight space and reading performance depend on factors that vary idiosyncratically. Together, these considerations make it unlikely that characterizing organization as the location in a generic weight space would have much utility.

⁹Although not explicitly demonstrated, this pattern is suggested by the results of simulation of the triangle model investigating the impact of semantics on the operation of the orthographic phonological pathway (e.g., Plaut et al., 1996; Harm & Seidenberg, 2004; Dilkina et al. 2008; Welbourne et al., 2011).

assumptions), the organization of individual B is characterized by greater reliance on the lexical/semantic pathway and less reliance on large grain-size regularities.

Finally, the contrast between individuals A and C illustrates how cross-language differences are captured in this framework. Note that in the control space A and C only differ on an exogenous dimension—which might correspond, say, to the ‘orthographic depth’ of the writing system each individual learns to read. Assume that C’s writing system is shallower than A’s (upward in Figure 1). As a consequence, relative to A, the organization of C’s reading system is less reliant on semantics and less reliant on larger grain sizes.

It is important to note that even though it captures a number of evidence-based conjectures, Figure 1 is purely hypothetical. It is merely meant to illustrate the general structure of the theoretical framework being proposed, although it also serves to highlight some of the challenges that a theory developed within this framework must confront. First, the dimensions of the organizational space must be identified. One possibility is that the dimensions correspond to a battery of experimental and standardized measures (depicted as the components of the reading profiles in Figure 1.). This approach is appealing in that it would make it relatively easy to determine an individual reader’s position in organizational space. On the other hand, it would make the theory more descriptive than explanatory and comparisons across languages or skill levels would be challenging. A more promising possibility is to identify theoretically motivated dimensions, such as the division of labor (see above), the sensitivity to orthographic units of different grain sizes (e.g. Ziegler & Goswami, 2005), or lexical quality (Perfetti, 2007). Alternatively, it may be possible to align the dimensions of the organization space with the degree to which the reader is attuned to the various statistical properties of the writing system.

A second challenge is to identify the determinants of the space itself: What determines which organizations are even possible? Presumably, there are two sets of constraints: biologically determined constraints reflecting the properties of the perceptual, learning, memory, and perhaps even languages-specific neurocomputational systems engaged during reading, and constraints that stem from the task demands of reading and the characteristics of the writing system. Many of these constraints would depend on the value of control parameters of the sort discussed above.

A third challenge is to ground the theory in neurobiological data. One aspect of this challenge is to identify the mapping (assuming one exists) between the components of the computational model and neuroanatomical structures (Rueckl & Seidenberg, 2009; Taylor, Rastle, & Davis, 2012). Indeed, it is not clear that there is a one-to-one mapping between organizations at the neural and computational levels—possibly, a given computational organization can be instantiated in the brain in more than one way.

A fourth challenge is to elucidate the nature of the learning mechanism(s) that underlie changes in organization. For example, to what extent does reading acquisition draw on domain-general statistical learning processes (Frost, 2012, Frost et al., 2015)? Does the organization of the reading system depend on consolidation processes (Davis & Gaskell,

2009)? What role do subcortical structures such as the thalamus play in reading acquisition (Pugh et al., 2013)?

Finally, a more empirical challenge is to develop the data sets and methodological approaches that are necessary for testing theories couched at the organizational level. To a large extent, the data that have been used to develop and test current theories of word reading were generated by experiments geared towards identifying the central tendencies of a population of readers. In such experiments, variability is little more than a nuisance. In the approach advocated here, variability is the signal rather than the noise, and few of the (thousands of) word reading studies that have been published to date are well-suited for testing hypotheses specifically concerned with this variability (e.g., that reading organization converges (or diverges) over the course of learning).

Summary

Theories of word reading have evolved over the years. Qualitative box-and-arrow models gave rise to more explicit, nuanced, and sophisticated computational models. Often, these computational models have ignored learning, focusing on whether a static (and stipulated) organization could account for the kinds of behavior observed in experimental word recognition tasks. Increasingly, however, computational models have incorporated a learning mechanism, and thus address both how words are read and, at a slower time scale, how the organization of the reading system changes with experience. The stage is now set to develop theories that explicitly address how and why the organization of the reading system varies. Indeed, although individual differences have long been of concern to research on reading acquisition and reading disability, the emerging interest in the differences among adult readers, both across (e.g., Frost, 2012; Share, 2008) and within (e.g., Andrews, 2012; Kuperman & Van Dyke, 2013; Welcome, & Joanisse, 2012) linguistic communities, is rather remarkable.

In addition to a focus on learning, the development of such theories will require a greater emphasis on the relationship between a model's control parameters and behavior. Interestingly, even for theories focused on central tendencies, methods for systematically characterizing this relationship are starting to emerge (e.g., Pitt, Kim, Navarro, & Myung, 2006), and it is worth noting that on occasion models of reading have been fitted to individual subject data by assuming that individual differences reflect underlying differences in a model's control parameters (Yap et al., 2012; Ziegler et al., 2008). However, it is also important to note that in a system that learns the relationship between its control parameters and its behavior may be variable, and this variability can arise from factors other than intrinsic or measurement noise. This suggests that a mediating construct is needed to capture the relationship between a system's control parameters and its behavior. In the discussion above, the notion of an *organizational space* was proposed as a candidate for this role and some of the conceptual and empirical challenges for such a theory were identified.

Acknowledgments

This research was supported by NIH grant P01 HD001994 to Jay Rueckl (PI). The author wishes to thank Ken Pugh, Ram Frost, Mark Seidenberg, and Jason Zevin for their influence on the ideas presented in this paper.

References

- Andrews, S. Individual differences in skilled visual word recognition and reading. In: Adelman, James, editor. *Visual Word Recognition Volume 2: Meaning and context, individuals and development*. Psychology Press; Sussex, UK: 2012. p. 151-172.
- Andrews S, Scarratt DR. Rule and analogy mechanisms in reading nonwords: Hough dou peapel rede gnew wirds? *Journal of Experimental Psychology: Human Perception & Performance*. 1998; 24:1052–1088.
- Backman J, Bruck M, Herbert M, Seidenberg M. Acquisition and use of spelling-sound information in reading. *Journal of Experimental Child Psychology*. 1984; 38:114–133.
- Coltheart M. Acquired dyslexias and the computational modeling of reading. *Cognitive Neuropsychology*. 2006; 23:96–109. [PubMed: 21049323]
- Coltheart M, Rastle K, Perry C, Langdon R, Ziegler J. DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*. 2001; 108:204–256. [PubMed: 11212628]
- Davis MH, Gaskell MG. A complementary systems account of word learning: neural and behavioural evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2009; 364(1536): 3773–3800.
- Dilkina K, McClelland JL, Plaut DC. A single-system account of semantic and lexical deficits in five semantic dementia patients. *Cognitive Neuropsychology*. 2008; 25:136–164. [PubMed: 18568816]
- Elman, JL.; Bates, EA.; Johnson, MH.; Karmiloff-Smith, A.; Parisi, D.; Plunkett, K. *Rethinking Innateness: A Connectionist Perspective on Development*. MIT Press; Cambridge, MA: 1996.
- Frost R. Toward a strong phonological theory of visual word recognition: True issues and false trails. *Psychological Bulletin*. 1998; 123:71–99. [PubMed: 9461854]
- Frost R. Towards a universal model of reading. *Behavioral and Brain Sciences*. 2012; 35:263–279. [PubMed: 22929057]
- Frost R, Armstrong BC, Siegelman N, Christiansen MH. Domain generality versus modality specificity: The paradox of statistical learning. *Trends in Cognitive Sciences*. 2015; 19:117–125. [PubMed: 25631249]
- Frost R, Katz L, Bentin S. Strategies for visual word recognition and orthographical depth: a multilingual comparison. *Journal of Experimental Psychology: Human Perception & Performance*. 1987; 13:104–115. [PubMed: 2951484]
- Grainger J, Jacobs AM. Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*. 1996; 103:518–565. [PubMed: 8759046]
- Griffiths TL, Chater N, Kemp C, Perfors A, Tenenbaum JB. Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences*. 2010; 14(8):357–364. [PubMed: 20576465]
- Harm MW, McCandliss BD, Seidenberg MS. Modeling the successes and failures of interventions for disabled readers. *Scientific Studies of Reading*. 2003; 7:155–182.
- Harm MW, Seidenberg MS. Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*. 1999; 106:491–528. [PubMed: 10467896]
- Harm MW, Seidenberg MS. Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*. 2004; 111:662–720. [PubMed: 15250780]
- Holyk GG, Pexman PM. The elusive nature of early phonological priming effects: Are there individual differences? *Brain and Language*. 2004; 90:353–367. doi:10.1016/S0093-934X(03)00447-4. [PubMed: 15172552]
- Kuperman V, Van Dyke JA. Reassessing word frequency as a determinant of word recognition for skilled and unskilled readers. *Journal of Experimental Psychology: Human Perception and Performance*. 2013; 39(3):802–823. doi:10.1037/a0030859. [PubMed: 23339352]
- Landi N, Frost SJ, Mencl WE, Preston JL, Jacobsen LK, Lee M, Yrigollen C, Pugh KR, Grigorenko EL. The COMT Val/Met polymorphism associated with reading-related skills and consistent patterns of functional neural activation. *Developmental Science*. 2013; 16:13–23. [PubMed: 23278923]

- Lerner I, Armstrong BC, Frost R. What can we learn from learning models about sensitivity to letter-order in visual word recognition? *Journal of Memory and Language*. 2014; 77(C):1–19.
- Manis F, Seidenberg M, Doi L, McBride-Chang C, Peterson A. On the basis of two subtypes of developmental dyslexia. *Cognition*. 1996; 58:157–195. [PubMed: 8820386]
- McClelland JL, Rumelhart DE. An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*. 1981; 88:375–407.
- McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*. 1995; 102(3):419–457. [PubMed: 7624455]
- Morton J. Interaction of information in word recognition. *Psychological Review*. 1969; 76:165.
- Norris D. The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*. 2006; 113:327–357. [PubMed: 16637764]
- Paulesu E, McCrory E, Fazio F, Menoncello L, Brunswick N, Cappa S, Cotelli M, Cossu G, Corte F, Lorusso M, Pesenti S, Gallagher A, Perani D, Price C, Frith C, Frith U. A cultural effect on brain function. *Nature Neuroscience*. 2000; 3:91–96. [PubMed: 10607401]
- Pitt MA, Kim W, Navarro DJ, Myung JI. Global model analysis by parameter space partitioning. *Psychological Review*. 2006; 113:57–83. [PubMed: 16478301]
- Plaut DC. Computational modeling of word reading, acquired dyslexia, and remediation. *Converging Methods in Reading and Dyslexia*. 1999:339–372.
- Plaut DC, Booth JR. Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*. 2000; 107:786–823. [PubMed: 11089407]
- Plaut DC, Gonnerman LM. Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*. 2000; 15:445–485.
- Plaut DC, McClelland JL, Seidenberg MS, Patterson K. Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*. 1996; 103:56–115. [PubMed: 8650300]
- Perfetti C. Reading ability: Lexical quality to comprehension. *Scientific Studies of Reading*. 2007; 11:357–383.
- Perry C, Ziegler JC, Zorzi. Beyond single syllables: Large-scale modelling of reading aloud with the connectionist dual process (CDP++) model. *Cognitive Psychology*. 2010; 61:2, 106–151.
- Perry C, Ziegler JC, Zorzi M. When silent letters say more than a thousand words: An implementation and evaluation of CDP++ in French. *Journal of Memory and Language*. 2014; 72:98–115. DOI: 10.1016/j.jml.2014.01.003.
- Pritchard SC, Coltheart M, Palethorpe S, Castles A. Nonword reading: Comparing dual-route cascaded and connectionist dual-process models with human data. *Journal of Experimental Psychology: Human Perception and Performance*. 2012; 38:1268–1288. [PubMed: 22309087]
- Pugh KR, Landi N, Preston JL, Mencl WE, Austin AC, Sibley DE, et al. The relationship between phonological and auditory processing and brain organization in beginning readers. *Brain and Language*. 2013; 125(2):173–183. [PubMed: 22572517]
- Pugh KR, Frost SJ, Rothman DL, Hoefl F, Del Tufo SN, Mason GF, et al. Glutamate and Choline Levels Predict Individual Differences in Reading Ability in Emergent Readers. *Journal of Neuroscience*. 2014; 34(11):4082–4089. doi:10.1523/JNEUROSCI.3907-13.2014. [PubMed: 24623786]
- Rueckl JG. The dynamics of visual word recognition. *Ecological Psychology*. 2002; 14(1–2):5–19.
- Rueckl JG. The limitations of the reverse-engineering approach to cognitive modelling. *Behavioral and Brain Sciences*. 2012; 35(5):305. [PubMed: 22929543]
- Rueckl, JG.; Seidenberg, MS. Computational modeling and the neural bases of reading and reading disorders. In: Pugh, KR.; McCardle, P., editors. *How children learn to read: Current issues and new directions in the integration of cognition, neurobiology and genetics of reading and dyslexia research and practice*. 1st ed.. Psychology Press; London (England): 2009. p. 99-131.
- Rumelhart, DE.; McClelland, J.; the PDP Research Group. , editors. *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 1: Foundation. MIT Press; Cambridge, MA:

- Rumelhart, DE.; Hinton, G.; Williams, R. Learning internal representations by error propagation. In: Rumelhart, DE.; McClelland, J.; the PDP Research Group. , editors. *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 1: Foundation. MIT Press; Cambridge, MA: 1986. p. 318-362.
- Seidenberg MS. Connectionist Models of Word Reading. *Current Directions in Psychological Science*. 2005; 14(5):238–242.
- Seidenberg, MS. Orthography and the Brain-Gene-Behavior Link. 2011. Reading in different writing systems: One architecture, multiple solutions. P. McCardle, B. Miller, J. Lee & O. Tzeng. *Dyslexia across languages*; p. 151-174.
- Seidenberg MS, McClelland JL. A distributed, developmental model of word recognition and naming. *Psychological Review*. 1989; 96:523–568. [PubMed: 2798649]
- Seidenberg MS, Plaut DC, Petersen AS, McClelland JL, McRae K. Nonword pronunciation and models of word recognition. *Journal of Experimental Psychology: Human Perception and Performance*. 1994; 20:1177–1196. [PubMed: 7844510]
- Share DL. On the Anglocentricities of Current Reading Research and Practice: The Perils of Overreliance on an “;Outlier” Orthography. *Psychological Bulletin*. 2008; 134:584–615. doi: 10.1037/0033-2909.134.4.584. [PubMed: 18605821]
- Strain E, Herdman CM. Imageability effects in word naming: an individual differences analysis. *Canadian Journal of Experimental Psychology*. 1999; 53:347–359. [PubMed: 10646206]
- Taylor JSH, Rastle K, Davis MH. Can Cognitive Models Explain Brain Activation During Word and Pseudoword Reading? A Meta-Analysis of 36 Neuroimaging Studies. *Psychological Bulletin*. 2012; 139(4):766–791. [PubMed: 23046391]
- Treiman R, Kessler B, Bick S. Influence of consonantal context on the pronunciation of vowels: A comparison of human readers and computational models. *Cognition*. 2003; 88:49–78. [PubMed: 12711153]
- Ullman MT. Contributions of memory circuits to language: The declarative/procedural model. *Cognition*. 2004; 92(1):231–270. [PubMed: 15037131]
- Van Orden GC, Holden JG, Podgornik MN, Aitchison CS. What swimming says about reading: Coordination, context, and homophone errors. *Ecological psychology*. 1999; 11(1):45–79.
- Welbourne SR, Woollams AM, Crisp J, Lambon Ralph MA. The role of plasticity-related functional reorganization in the explanation of central dyslexias. *Cognitive Neuropsychology*. 2011; 28:65–108. [PubMed: 22122115]
- Welcome SE, Joanisse MF. Individual differences in skilled adult readers reveal dissociable patterns of neural activity associated with component processes of reading. *Brain and Language*. 2012; 120(3):360–371. doi:10.1016/j.bandl.2011.12.011. [PubMed: 22281240]
- Woollams A, Lambon Ralph MA, Plaut DC, Patterson K. SD-squared: On the association between semantic dementia and surface dyslexia. *Psychological Review*. 2007; 114:316–339. [PubMed: 17500629]
- Yang J, Shu H, McCandliss BD, Zevin JD. Orthographic influences on division of labor in learning to read Chinese and English: Insights from computational modeling. *Bilingualism: Language and Cognition*. 2012; 16(02):354–366. doi:10.1017/S1366728912000296.
- Yap MJ, Balota DA, Sibley DE, Ratcliff R. Individual differences in visual word recognitions: Insights from the English Lexicon Project. *Journal of Experimental Psychology: Human Perception and Performance*. 2012; 38:53–79. [PubMed: 21728459]
- Zevin JD, Seidenberg MS. Consistency effects and individual differences in nonword naming: A comparison of current models. *Journal of Memory and Language*. 2006; 54:145–160.
- Ziegler JC, Goswami U. Reading acquisition, developmental dyslexia, and skilled reading across languages: A psycholinguistic grain size theory. *Psychological Bulletin*. 2005; 131:3–29. [PubMed: 15631549]
- Ziegler JC, Bertrand D, Tóth D, Csépe V, Reis A, Fásca L, et al. Orthographic depth and its impact on universal predictors of reading: A cross-language investigation. *Psychological Science*. 2010; 21:551–559. [PubMed: 20424101]

- Ziegler JC, Castel C, Pech-Georgel C, George F, Alario F-X, Perry C. Developmental dyslexia and the dual route model of reading: Simulating individual differences and subtypes. *Cognition*. 2008; 107:151–178. [PubMed: 17959161]
- Ziegler JC, Perry C, Coltheart M. The DRC model of visual word recognition and reading aloud: An extension to German. *European Journal of Cognitive Psychology*. 2000; 12:413–430.
- Ziegler JC, Perry C, Zorzi M. Modelling reading development through phonological decoding and self-teaching: Implications for dyslexia. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2014; 369(1634) art. no. 20120397. DOI: 10.1098/rstb.2012.0397.




Figure 1.
A depiction of the proposed theoretical framework. (See text for explanation.)