Published in final edited form as:

J Mem Lang. 2009 August 2; 61(2): 238–257. doi:10.1016/j.jml.2009.05.001.

Simulating Language-specific and Language-general Effects in a Statistical Learning Model of Chinese Reading

Jianfeng Yang^{1,2}, Bruce D. McCandliss¹, Hua Shu², and Jason D. Zevin^{1,*}

¹Sackler Institute, Weill Cornell Medical College, USA

²State Key Laboratory of Cognitive Neuroscience and Learning, Beijing Normal University, China

Abstract

Many theoretical models of reading assume that different writing systems require different processing assumptions. For example, it is often claimed that print-to-sound mappings in Chinese are not represented or processed sub-lexically. We present a connectionist model that learns the print to sound mappings of Chinese characters using the same functional architecture and learning rules that have been applied to English. The model predicts an interaction between item frequency and print-to-sound consistency analogous to what has been found for English, as well as a language-specific regularity effect particular to Chinese. Behavioral naming experiments using the same test items as the model confirmed these predictions. Corpus properties and the analyses of internal representations that evolved over training revealed that the model was able to capitalize on information in “phonetic components” – sub-lexical structures of variable size that convey probabilistic information about pronunciation. The results suggest that adult reading performance across very different writing systems may be explained as the result of applying the same learning mechanisms to the particular input statistics of writing systems shaped by both culture and the exigencies of communicating spoken language in a visual medium.

Keywords

Reading; Computer Simulation; Cross Cultural Differences

Over the past three decades, computational models have become increasingly sophisticated in accounting for a broad range of phenomena and specifying the mechanisms underlying skilled reading and its acquisition (see reviews in, e.g., Rayner, Foorman, Perfetti, Pesetsky, & Seidenberg, 2001; Plaut, 2005). The vast majority of this work has been done in English, and has thus focused on issues arising from the particularities of its writing system (Share, 2008). This has led to the construction of models that implement relatively writing system specific assumptions, such as the inclusion of distinct processing mechanisms for “sub-lexical” and “lexical” translation from spelling to sound (Coltheart, Curtis, Atkins, & Haller, 1993; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001). An alternative approach has been to assume that reading skill is acquired by way of domain-general learning mechanisms that operate over distributed representations of basic levels of information, such as orthography,

© 2009 Elsevier Inc. All rights reserved.

Address for Correspondence: Jason D. Zevin, Sackler Institute, Weill Cornell Medical College, 1300 York Ave., Box 140, New York, NY, 10065, Tel: 212-746-3781, E-mail: jdz2001@med.cornell.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

phonology and semantics (Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989). Here we present a general connectionist model of Chinese print-to-sound translation that implements the same functional architecture and learning rules as models that have been previously applied to English (Harm & Seidenberg, 1999, 2004; Treiman, Kessler, Zevin, Bick, & Davis, 2006; Zevin & Seidenberg, 2002, 2006). The model provides a computationally explicit theoretical account of the role of sublexical phonology and the emergence of functional units in Chinese reading.

“Orthographic Depth” and print-to-sound translation in Chinese

The notion of “orthographic depth” (Bentin & Frost, 1987) provides a descriptive framework that captures important differences among writing systems. In “shallow” orthographies, each character or group of characters corresponds with a high degree of consistency to a single speech sound. An extreme example of this is the Hangul system in Korean, in which each syllable is transcribed as a set of one to four *jamo*, each of which in turn comprises a set of strokes that each indicate a specific phonetic feature, assuring that words with the same pronunciation are identical to one another in the script, and that words with similar pronunciations are written similarly (Lee, 2000). Although English spelling is famously complicated (Malone, 1925; Venezky, 1999), at its core is an alphabet of letters that correspond roughly to individual speech sounds (Venezky, 1970) which can be described in terms of sublexical mapping “rules”. In the dual-route framework, these rules are supplemented by a lexical route that contains the correct pronunciations for known words (Coltheart et al., 2001; Zorzi, Houghton, & Butterworth, 1998).

Chinese, in contrast, is an example of an extremely “deep” orthography (DeFrancis, 1989) in that the pronunciation of a character cannot be computed sound-by-sound from its constituent parts. Chinese orthography has a roughly hierarchical organization, with five types of strokes combined to form *radicals*, orthographic units that may carry probabilistic information about meaning and sound. Unlike letters, radicals do not contain componential information about pronunciation (Mattingly, 1987). There is no relationship between, e.g., the first radical in a character and the first phoneme in the spoken word it represents.

Further, radicals are often organized into larger units such as phonetic and semantic components. The majority (85%) of characters in Chinese are phonograms (Zhu, 1988), which consist of a semantic component that provides information about the meaning of the character, and a phonetic component that provides information about the character’s pronunciation (Li & Kang, 1993). Although the pronunciation of each character can be probabilistically determined by its phonetic component, it is sometimes entirely arbitrary, so that very similarly written words often have completely different pronunciations, and many homophones share no orthographic features.

The effect of orthographic depth on the organization of the reading system is often described in terms of a bias toward using one or another of the component processes hypothesized to be necessary for English (Frost, Katz, & Bentin, 1987). Visual word recognition in shallow orthographies is thought to involve relatively limited engagement of lexical processes, because sub-lexical translation from spelling to sound is so automatic and consistently accurate. For example Korean Hangul characters are thought to be read predominately via sub-lexical processing unless task demands create a strong bias toward lexical processing (Kang & Simpson, 2001, also see Raman, Baluch, & Besner, 2004 for Turkish; Havelka, & Rastle, 2005 for Serbian).

Many existing models of Chinese character reading reflect the converse assumption, i.e., that spelling-to-sound translation is accomplished entirely via a “lexical” mechanism (Perfetti, Liu, & Tan, 2005; Zhou & Marslen-Wilson, 1999a). The lexical mechanism for print to sound

translation is most often described as relying on hierarchically organized levels of representation, with localist representations at each level, in the vein of the Interactive Activation model of English word recognition (McClelland & Rumelhart, 1981).

The architecture of the lexical mechanism in these models also reflects a number of assumptions about the organization of the reading system that are fairly specific to Chinese. For example, Taft and colleagues (Taft, 2006; Taft & Zhu, 1997; Taft, Zhu, & Peng, 1999) have proposed a multilevel interactive-activation model of word recognition in Chinese that assumes characters, radicals and strokes each have their own separate level of representation, and that these are organized serially. The assumption that radicals form a basic orthographic unit is well supported by a number of priming studies (Ding, Peng & Taft, 2004), but as we shall argue below, additional considerations suggest that radicals are not necessarily apt functional units for print-to-sound conversion.

Rethinking the Assumptions of Chinese Characters Reading

A number of empirical findings appear to conflict with the notion that print-to-sound translation is purely lexical in Chinese. First, some properties of the writing system suggest that a sub-lexical mechanism is plausible after all. Phonetic components contain probabilistic sub-lexical information about how the whole characters are pronounced. There is a great deal of evidence that this sub-lexical information is involved in reading aloud. When a phonogram's pronunciation matches its phonetic component exactly, it is called "regular" (Seidenberg, 1985; Shu & Zhang, 1987). Large effects of regularity — regular characters are named faster and more accurately — have been found in studies of both adults (Fang, Horng, & Tzeng, 1986; Hue, 1992; Lee, Tsai, Su, Tzeng, & Hung, 2005; Peng & Yang, 1997) and children (Shu, Anderson, & Wu, 2000; Shu & Wu, 2006; Xing, Shu, & Li, 2004; Yang & Peng, 1997). Although regularity effects are *prima facie* evidence for sub-lexical mappings from print-to-sound, they are also potentially consistent with models that assume purely lexical activation of phonology. It could be possible, for example, to construct a model such that when a character that has a lexical representation is presented — even as part of another character — it becomes activated, thus leading to a regularity effect.

Phonetic components also vary with respect to how reliably the words they appear in match their canonical pronunciation, so that even items that are "regular", as described above, can vary in "consistency" in a manner directly analogous to English (Jared, McRae, & Seidenberg, 1990). Just as a completely regular pronunciation of a rime in English can be difficult to compute if there are many irregular neighbors (e.g., DOLL vs. ROLL, POLL, TOLL), a completely regular character can be difficult to pronounce if many characters containing the same phonetic component have irregular pronunciations (Fang et al., 1986; Hue, 1992). These effects argue more strongly for some form of sub-lexical print-to-sound translation, because they reflect distributional information related to shared sub-lexical units (phonetic components) among characters. Although it has been suggested that consistency effects in English might arise from activation of similar words in the lexical route, for example in the Dual-Route Cascade model (Coltheart et al., 2001), subsequent work with this model has shown that where this model correctly simulates consistency effects, it is due to sub-lexical processing (Seidenberg & Plaut, 2006). It is thus an open question whether a purely lexical model of print-to-sound translation for Chinese could correctly simulate consistency effects.

The role of phonetic components in print to sound translation also poses a problem for the notion of radicals as the basic unit of character reading, because phonetic components can comprise one or more individual radicals. Some single radicals can also function as phonetic components. For example, the radical “扌” is pronounced /pɑI/ when it appears as a word, and is also the phonetic component of a number of words that share the same pronunciation (e.g.,

“柏伯/pai/”). The same radicals may participate in many different phonetic components, however. For example, the “𠂇” radical also occurs in characters that do not share its pronunciation at all (e.g. “皇/hwɑŋ/”), which in turn serve as phonetic components of larger characters (皇煌煌煌: /hwɑŋ/). Complicating matters further, the same simple radical occurs in many contexts (𠂇/sau/, 𠂇/xau/) without any relation to the whole character’s pronunciation. Thus, although radicals clearly function as a unit at the level of orthographic organization, (Taft & Zhu, 1997, Zhang, Perfetti & Yang, 1999, Taft, 2006, Ding, Peng & Taft, 2004), they are not usually the most relevant functional unit for print-to-sound translation. This raises an important challenge for models of reading in Chinese-- how can models explain the emergence of phonetic components in print-to-sound mappings when they are not readily identifiable as “basic units” from the surface features of the orthography?

A General Framework across Orthographic Depth from a Connectionist Perspective

Many theorists assume that different scripts necessitate different processing assumptions. For example, in discussing the possibility of extending their model of English reading to other languages, Coltheart et al. (2001) assert that “the Chinese, Japanese, and Korean writing systems are structurally so different from the English writing system that a model like the DRC model would simply not be applicable.” (p. 236). An alternative approach is to assume that the basic processes underlying reading are essentially the same across writing systems, and that differences in how the reading system becomes organized are due to the influence of statistical learning mechanisms operating over distributed orthographic, phonological and semantic representations (Seidenberg, 1992; Seidenberg & McClelland, 1989). Thus, rather than assuming that Chinese and English require entirely different processing mechanisms, this approach seeks to explain processing differences between the two languages in terms of properties of the orthography that lead to different learning outcomes. The most important of these in the current case are the grain size (Ziegler & Goswami, 2005) and degree of arbitrariness (Seidenberg, 1992) of mappings across writing systems.

English can be described as having regularities at multiple grain sizes (Ziegler & Goswami, 2005). For example, each spoken consonant in the English word BREED is represented by a single letter, whereas the vowel is written as a two-letter combination. These written forms of single phonemes – or “graphemes” -- are all completely regular, in the sense that they are each assigned their most typical pronunciation in this word. In contrast, the word BREAD contains an atypical pronunciation for the grapheme EA. Words that contain such mappings are called “irregular” in some models (Coltheart et al., 2001). Note, however, that the spelling-to-sound mapping for BREAD is supported by many other examples, e.g., HEAD, THREAD, and TREAD. Statistical regularities at this level of detail give rise to consistency effects (Jared, McRae, & Seidenberg, 1990). Thus, subsyllabic mappings from spelling to sound are dominant in the English writing system, with multiple grain sizes contributing to the spelling-to-sound mapping, sometimes in conflicting ways.

Note that describing English in this manner reveals an important parallel between English and Chinese. In both systems, a basic orthographic unit (the letter in English, the radical in Chinese) can be recombined into larger functional units for print-to-sound translation. Indeed, regularities at multiple grain sizes are clearly used productively in generalization, as demonstrated by studies of nonword reading in adults (Treiman, Kessler, & Bick, 2002) and children (Treiman & Kessler, 2006)—effects that are accurately simulated in connectionist

¹Here and throughout, transcriptions are in International Phonetic Alphabet. Note that in the Pinyin romanization system, the voiceless, unaspirated stop /p/ is written as a “b.”

models that explicitly capitalize on multiple grain sizes simultaneously during learning (Treiman et al., 2006; Zevin & Seidenberg, 2006).

The key difference between English and Chinese, then, is that whereas functional print-to-sound units in English can exist at a number of different sizes, in Chinese, regularities exist almost entirely in the mapping of phonetic components to whole syllables (Leong, 1997; Mattingly, 1987). The difference in the grain size at which regularities exist drives a difference in the degree of arbitrariness between the two writing systems. English spelling, while highly inconsistent, is never entirely arbitrary. Even a very strange word such as YACHT has some predictability to it (the Y, A and T are assigned pronunciations common in other contexts). In Chinese this is not necessarily the case. The many pronunciations for characters containing the “*台/t'ai/*” component, for example, are a highly heterogeneous set (see Figure 1A)². Whereas in English, the spelling of an unfamiliar word provides partial cues to its pronunciation, in Chinese characters, the distribution of pronunciations for any given phonetic component may contain a number of highly dissimilar forms. Thus, the differences between the two writing systems can be described in terms of factors that have well-characterized effects on models that employ domain- and language-general statistical learning mechanisms.

Previous computational models of Chinese reading

The Lexical Constituency Model (LCM, Perfetti et al., 2005) is an implemented model based loosely on the interactive activation model of reading in English. The focus of research in the LCM framework is establishing a role for phonological processing in reading, despite what is described as a lack of reliable sub-lexical phonological information in the writing system. In the authors' view this reflects a general principle about reading across languages, i.e. that the natural medium for language is speech, and therefore reading necessarily involves accessing phonological representations. However, the processes by which print to sound translation is accomplished are thought to be defined by the language. Unlike alphabetic writing systems, which the same authors describe as having “assembled phonology as a sub-lexical mechanism” (Perfetti et al., 2005, p. 43) in addition to a lexical look-up mechanism (Coltheart et al., 1993), Chinese is described as depending on only the lexical look-up mechanism. The input to this model comprises localist representations of radicals. These feed forward to a lexicon of localist representations of whole characters, which in turn are connected to their pronunciations in a phonological layer. Representations of character meaning are activated jointly by input from the phonological and lexical orthographic layers. Critically, there is no direct input from the radical layer to the phonological layer (although many radicals are also characters in their own right, and have redundant representations in the lexical layer which are connected to their pronunciation). In this model, then, access to pronunciation is strictly lexical, in the sense that sub-character components cannot directly activate their pronunciations. Thus, although the modeling framework is fairly general, the implementation embodies a number of writing-system-specific assumptions about how the reading system is organized for Chinese.

In addition to the LCM (Perfetti et al., 2005), there have been a number of connectionist models of Chinese reading. These have tended to focus on language-specific phenomena. For example, Hsiao and Shillcock (2005, 2006) focused on the interaction between regularity and the position of phonetic and semantic radicals, but didn't explore frequency by consistency interactions. Regularity in this model was defined in terms of the canonical pronunciation of the phonetic

²Fully 48% of irregular words share neither onset nor rime with their regular counterparts, making them dramatically less predictable than even “strange” words in English. Interestingly, however, there is *some* subsyllabic regularity, such that 18% of irregular forms share their onset with regular family members, and 35% share a rime. These are substantially greater than would be predicted if pronunciations were uniformly distributed (spoken Standard Mandarin has 23 onsets and 39 rimes). Note, however, that due to the non-componential nature of the mapping, it is not possible to know whether a particular phonetic component predicts the onset or the rime (or the vowel, or the tone, which we have not explored).

radical This model also includes some interesting details about the potential role of hemispheric asymmetry in reading aloud that makes specific predictions about experiments in which stimuli are presented to either the left or right visual field. This model uses localist representations for phonetic and semantic components, and therefore can say little about the emergence of functional units at multiple grain sizes -- the functional unit of analysis for print-to-sound translation is presupposed by the selection of the training corpus, and pre-coded into the network's architecture.

Another model, developed by Xing, Shu & Li (2002, 2004) focused on simulating how children acquire reading skill in Chinese based on Self-Organized Feature Mapping (SOFM) between orthographic and phonological similarity. The SOFM model captured the development of regularity, consistency and their interaction with frequency, correctly simulating some aspects of children's acquisition of characters during the elementary school years. One limitation of this model was that training corpus was quite small (about 300 words), and it is unclear that the particular formalism employed would generalize to other training sets, or scale up to a larger lexicon. A more serious limitation of the model was the way regularity effects were simulated: Each time a word was presented, the pronunciations both of the whole character and its phonetic component were input simultaneously to the model. Thus, the model shows that statistical learning mechanisms can integrate sub-lexical and lexical information about pronunciation, but it assumes that the phonetic component is an a priori functional unit in print-to-sound mapping, and that its most frequent pronunciation is explicitly related to whole character's pronunciation each time it is encountered, instead of learning the mapping from print to sound.

Finally, there is a preliminary report of a model that employs a similar architecture to models that have been applied to English (Chen & Peng, 1994). This model has orthographic and phonological representations that do not build in any assumptions about the role of phonetic components in computing pronunciations from print, and yet it correctly simulates the interaction between frequency and regularity as observed in a number of contemporary studies (Fang et al., 1986, Hue, 1992). Although these are encouraging preliminary results, the model did not address consistency effects, nor were any analyses conducted on the acquisition of reading skill over the course of development. Thus, existing models of reading in Chinese have either been implemented with Chinese-specific assumptions about the processes (Perfetti et al., 2005), and functional units (Hsiao & Shillcock, 2005, 2006; Xing, Shu & Li, 2002, 2004), or have addressed a relatively limited range of phenomena (Chen & Peng, 1994).

The current model

Here we adapt a computational model that has been successful in explaining a range of English reading phenomena to simulating skilled reading aloud in Chinese. In this framework, aspects that are common to the two reading systems, such as the frequency by consistency interaction, can be accounted for in terms of the learning mechanisms that underlie acquisition: Frequency effects are simply practice effects on particular items, whereas consistency effects reflect the influence of statistical patterns across many similar items, and emerge from a generic tendency to preserve similarity: i.e., most statistical learning mechanisms are predisposed to map similar inputs onto similar outputs. Thus, the most difficult items both in acquisition and skilled performance are those with statistically rare print to sound correspondences that are also encountered infrequently (Lee et al., 2005; Jared, 2002). At the same time, the adaptation of these mechanisms to different properties of the writing systems (grain size, arbitrariness) explain language-specific effects as emerging from statistical patterns shaped by historical and linguistic forces.

The model presented here has the same functional architecture and learning rules that have been used in a number of English studies (Harm & Seidenberg, 1999; Treiman et al., 2006; Zevin & Seidenberg, 2002, 2006), but with input and output representations modified to represent Chinese orthography and phonology. In Study 1, we demonstrate that the model correctly simulates both the effects of regularity and consistency, and their interactions with frequency. In Study 2, we explore the frequency by consistency interaction in detail, in light of the statistical properties of the lexicon and potential biases in the stimulus set designed for Study 1. Here again, the model correctly predicts human performance in a word naming task. Finally, a series of analyses of the training corpus and the internal representations that support the model's performance -- including how these develop over time were conducted. These analyses provide new insights into how the phonetic component emerges as a functional unit in learning to read Chinese.

Study 1

Connectionist model

Model architecture—Every unit in the orthographic input layer was connected to the 200 hidden units, which in turn were connected to every unit in the phonological output layer (See Figure 2A). In addition, every unit in the output layer was connected back to the output layer (including auto-connections) both directly and via a set of 50 "cleanup" hidden units, forming an attractor structure (Hinton & Shallice, 1991; Plaut & Shallice, 1993). This architecture differs from models used in English (e.g., Treiman et al., 2006; Zevin, & Seidenberg, 2002, 2006) only in the number of units used in each layer, and the fact that the orthographic and phonological representations are based on Chinese instead of English. Prior to training, all connections in the model were randomized to weights between -0.1 and 0.1 with a mean of 0.0 and a Gaussian distribution of values over the network.

Phonological representation—The phonological representation was based on a phonetic description of Standard Mandarin (Huang & Liao, 2002) and comprised five slots, one for the onset (in Chinese, the onset unit, or “*shengmu*” contains only the initial consonant), three for the rime (the “*yunmu*”, which in this case includes any semivowels, the nuclear vowel and codas) and a fifth slot composed of four units to represent lexical tone (See Figure 2B). The representations were centered so that the nuclear vowel always occurred in the second slot of the rime.

Each phoneme was encoded as a set of abstract phonetic features using a distributed representation with graded similarity. This is equivalent to representations that have been used in a number of English models (Harm & Seidenberg, 2004; Treiman et al., 2006; Zevin & Seidenberg, 2006). The four slots used to represent phonemes comprised 3 groups: 1) 8 units encoded manner (e.g., stop, affricate, fricative, semivowel/liquid, vowel), 2) 6 units encoded place (e.g., dental, retroflex, palatal, velar), 3) 8 units encoded impressionistic vowel quality, with 1 unit to code retroflexing, 2 to represent backness, 3 for height and 2 for lip rounding. Thus, each of the four phoneme slots had 22 units, which, taken together with the 4 units for lexical tone, makes 92 output units. Note that more units were used in some cases than was mathematically sufficient to represent the number of possible values in a feature space. This was done to encode the relative similarity of different features, for example, stop is coded as “1 0 0”, affricate as “1 1 0” and fricative as “1 1 1” on the units that encode degree of closure in the “manner” group.

Orthographic representation—The orthographic layer of the model was organized into a set of 9 *slots*, each of which was further organized into a variable number of *groups* containing a variable number of *units* that could take on values of either 0 or 1 – comprising 270 binary units in all (See Figure 2C). Each character was represented by a unique pattern of activity

over these units, based on a linguistic description (Xing, Shu, & Li, 2004) of Chinese orthography, adapted into a distributed representation that included features for both the hierarchical structure of each character (overall shape, relative position of radicals) and the orthographic details of which radicals appear in the character, and a finer-grained description of the radicals in terms of strokes.

Two slots encoded the overall structure of each character. One slot, containing seven groups of between 2 and 4 units each, represented overall character shape, with each encoding one type: left-right, top-bottom, left-middle-right, top-middle-bottom, cross, round and single. All together, 27 different sub-types of character structure were encoded using 18 units. The second slot comprised a single group of three units for the number of radicals in the character (one to seven).

Another seven slots were used to represent the radicals in the character, with each slot corresponding to a single radical according the definition of 560 unique radicals in the *Chinese Character Component Standard of GB13000.1 Character Set for information Processing* (1997). Radicals are roughly analogous to letters in an alphabetic script, except that they can be arranged in a variety of configurations within a character and do not necessarily encode for phonology. The variety of positions in which radicals can appear poses a version of the "dispersion problem" – i.e., that because the input representations are position-specific, what is learned about a letter in one position does not transfer to the same letter in a different position -- that has been noted in English (Plaut et al., 1996). In Chinese, the problem is more complex because of the greater possible number of spatial arrangements a character can take. For example, phonetic components typically appear on the right in a character with an overall left-right structure, and on the bottom in top-bottom characters. In the model, radicals were aligned in slots to reflect this similarity, as shown in Table 1. For left-right characters, the radicals from the left component were left-aligned into slots 1 through 3, and the radicals from the right component were right-aligned into slots 5 through 7, with the same overall scheme for top-bottom characters (with the top component left-aligned and the bottom component right-aligned). For left-middle-right and top-middle-bottom characters, the "middle" radicals were central-aligned in slot 4. When a phonetic component appeared alone in a character, it was represented in its most frequent position according its family members. This encoding was adopted in order to represent the similarity of characters containing the same components. It maintains the overlap among characters with very different shapes that share a phonetic component, which is critical for simulating both consistency and regularity effects.

All radical slots comprised at least 9 groups, three of which encoded overall structure, five of which encoded the "strokes" used to write the character, and a final group of 6 units used to disambiguate the 53 remaining radicals that could not be uniquely identified based on the first two groups. For six slots, a tenth group of 3 units encoded the relationship between the radical in that slot and adjacent radicals in terms of six different possible relations (e.g., left-right, surrounding, single). Groups for overall structure comprised 1) a group of 4 units to encode ten types of radical structure, -- e.g., left-right (ㄣ), top-bottom (㇏), surrounding (㇏), 2) a group of 3 units to encode six possible stroke relations -- e.g., crossing (乂, 十), separate (ㄣ, ㇏), connecting (ㄣ, ㇏), and 3) a group of 4 units to encode eleven possible modal positions. Five groups encoded for strokes, including 4 units to encode number of strokes from 1–10, and 3 units to encode five possible stroke types (horizontal, vertical, slanted, pointed and crooked) for the first, second, third and last stroke in the radical, respectively.

Training—A set of 4468 items from the Modern Chinese Frequency Dictionary (1986) was used to train the model. During training, the probability of using any character on a given trial was proportional to the square root of its frequency (Plaut et al., 1996), with raw frequencies

capped at 1,000. This ensured that low-frequency characters would be selected a reasonable number of times over the 3 million training trials.

Following Harm & Seidenberg (1999), we first pre-trained the phonological attractor model, and then trained the full reading model on the mapping from orthography to phonology. The continuous recurrent back-propagation algorithm (Pearlmutter, 1995) was used, with online learning, a learning rate of 0.005 and momentum of 0.9. On each trial, a character was selected and the orthographic units were clamped with the pattern corresponding to the writing of the character for 12 time ticks. Error computed at the phonological layer was computed after the first time tick activation was propagated forward to the output layer from time ticks 5–12 and a gradient based on this error was back propagated to update the connection weights.

Testing—A set of 120 items was chosen to test the frequency, consistency and regularity effects from the model on human participants: 20 from each cell resulting when three levels of regularity/consistency (regular-consistent, R-C; regular-inconsistent, R-I; irregular-inconsistent, I-I) were crossed with two levels of frequency (high, HF; low, LF). The mean frequency of occurrence was 475.35 per million for HF characters and 12.56 per million for LF characters, yielding a frequency manipulation of equivalent size across all three levels of regularity/consistency. Regular-consistent items contain phonetic components that appear only in characters that share the same onset and rime, regardless of tone. By definition, regular inconsistent items have the same pronunciation as their phonetic component, although other characters sharing the same phonetic component are pronounced differently. Conversely, irregular inconsistent items are pronounced differently from their phonetic components. The number of characters that share a phonetic component is called "family size," a property that was matched across the various frequency and regularity/consistency conditions, with the exception that the R-C items have a smaller family size than the others. Note that throughout the lexicon, R-C items have a much smaller number of neighbors than the other stimulus types; because larger family size have a facilitative effect on early visual recognition for naming R-C items (Feldman & Siok, 1999; Hsu, Lee & Tzeng, 2009), we matched the family size for inconsistent items only and included R-C items with a smaller family size than R-I or I-I items. Consistency is defined by the ratio of the frequency and number of "friends" (characters with the same pronunciation) and "enemies" (characters with different pronunciations, Peng & Yang, 1997). We therefore matched frequency weighted consistency value for inconsistent items. The consistency value is equal to summed frequency of friends divided by summed frequency of all family members (including the character itself, as in Shu, Chen, Anderson, Wu, & Xuan, 2003). Characters were also matched across all conditions for structure type, typical position of phonetic components, number of strokes and number of radicals.

Testing was carried out by presenting each character to the model, i.e., clamping the appropriate pattern on the input layer. We allowed the model to compute a pronunciation over 12 time ticks and the last tick was counted as the final output. Naming accuracy and sum squared error (SSE) were computed to test the model's performance. Accuracy was determined by applying a winner-take-all scoring system: for each slot on the output layer, we determined which phoneme was closest to the pattern on the output at the final time tick and reported this as the model's pronunciation; responses that did not match the correct pronunciation were scored as errors. SSE, a stand-in for response latency, was computed from the model's output at the second to last time tick by adding together the square of the difference between the model's output and the target for each unit. Because the model was run 40 times with different starting weights and random seeds, we were able to conduct both "by subjects" analyses that treated each run as a subject, and analyses with items as a random factor.

Results—At the end of training, 95% of the test items were named correctly. A 2 (high vs. low frequency) $\times$ 3 (type: R-C vs. R-I vs. I-I) ANOVA was conducted to examine the effects

of frequency, regularity and consistency, and their interaction (See Table 3). The overall correlation between SSE and response latency for all items in the behavioral experiment was 0.617 ($p < 0.01$).

This analysis revealed a significant frequency effect for both naming accuracy and SSE. The main effect of consistency/regularity was also significant for naming accuracy and SSE. A strong interaction between frequency and consistency/regularity was found for both accuracy and SSE, such that the effect of print-to-sound regularity and consistency were both larger for low than for high-frequency words. Performance was perfect for high-frequency regular words, whereas high-frequency irregulars elicited 1% errors. For low-frequency items, a strong effect of the consistency/regularity factor was observed in naming accuracy and SSE.

Large effects of stimulus type were observed in oneway ANOVAs for each level of frequency (Table 3). To further explore the effects of regularity and consistency for low-frequency items independently, Tukey tests (one-tailed) were conducted for both SSE and accuracy. The effect of consistency (R-I vs. R-C) was reliable for SSE $t_1(39) = 12.26, p < 0.01$; $t_2(38) = 2.29, p < 0.05$, but reliable only by participants for accuracy, $t_1(39) = 6.83, p < 0.01$; $t_2(38) = 1.68, p = 0.11$. A similar pattern was observed for regularity effects (R-I vs. I-I) among inconsistent words, with regularity effects reliable for SSE, $t_1(39) = 9.29, p < 0.01$; $t_2(38) = 4.57, p < 0.05$, and reliable by participants only for accuracy, $t_1(39) = 4.65, p < 0.01$; $t_2(38) = 1.65, p = 0.12$.

Behavioral experiment

Methods—To test predictions of the model, a character naming experiment was conducted with 39 undergraduates (19 male, 20 female) from Beijing Normal University. All participants were native speakers of Mandarin Chinese with normal or corrected-to-normal vision, aged between 17 and 25. They provided written informed consent and were paid for their participation. Participants sat a comfortable distance from the screen (about 60cm) and were instructed to read aloud single characters into a microphone as quickly and accurately as possible. On each trial, a fixation cross appeared for 500ms, after which the screen was cleared for 120ms and a single character was presented for up to 2000ms (or until a response was made). Stimuli were presented centrally, in white against a black background using 28pt *Songti* font. Stimulus presentation and response latency collection was controlled using DMASTR software (Forster & Forster, 2003). Stimuli were presented in a different, randomized order for each participant.

Results—Three participants' data were removed from the analysis because of error rates above 25%. Overall naming accuracy for the remaining 36 participants was 92.8%. Response latency and accuracy for each condition are summarized in figure 3, revealing effects that confirm each of the model's predictions.

A 2 (high vs. low frequency) $\times$ 3 (R-C vs. R-I vs. I-I) ANOVA was conducted both for naming latency and accuracy (See Table 4). The main effect of frequency was observed both for naming accuracy and response latency. The high frequency items were named faster and more accurately than low frequent items. Consistency/regularity influenced participants both on accuracy and response latency. Most critically, the predicted interaction between frequency and the consistency/regularity factor was observed for both accuracy and response latency.

For low frequency items, the effect of consistency (R-I vs. R-C) was reliable for both response latency, $t_1(35) = 6.55, p < 0.01$; $t_2(38) = 2.75, p < 0.01$, and for accuracy, $t_1(35) = 4.08, p < 0.01$; $t_2(38) = 2.37, p < 0.05$. A regularity effect (R-I vs. I-I) was also found for these items on response latency, $t_1(35) = 6.54, p < 0.01$; $t_2(38) = 3.56, p < 0.01$, and accuracy, $t_1(35) = 5.39, p < 0.01$; $t_2(38) = 3.05, p < 0.01$.

Discussion

A computational model of reading aloud in Chinese based on connectionist principles correctly simulated effects of regularity, consistency, frequency and their interaction, as observed in a behavioral study of reading aloud. In both the human and simulation data, regularity and consistency effects were each larger for low- than for high-frequency words. A regularity effect was found for inconsistent items, and was correctly simulated, demonstrating that regularity effects do not depend on categorical rules, but can emerge from the statistics of the writing system.

The interaction between frequency and consistency observed in Study 1 was quite large, and no consistency effect was found for high-frequency items. While this is broadly consistent with many studies of reading in English (Taraban & McClelland, 1987; Balota & Ferraro, 1993; Treiman, Mullennix, Bijeljac-Babic, & Richmond-Welty, 1995), the current results are in apparent conflict with some findings in the literature, particularly with respect to the strength of the frequency by consistency interaction and the lack of a consistency effect for high frequency items. In a series of experiments focusing on the neighborhood statistics that give rise to consistency effects, Jared (1997) found a relatively weak frequency by consistency interaction when neighborhood properties were tightly controlled between high- and low-frequency items. Although there was a trend toward an frequency by consistency interaction for both response latency (non-significant) and accuracy (significant by subjects, but not by items), this failure to replicate the “standard” effect under tight controls raised questions about the real basis of this interaction, which may be relevant to the current study.

The focus of the Jared (1997) study was on the consistency effect for high-frequency items. In a series of experiments, neighborhood properties (described in terms of “friends” and “enemies” as above) were shown to have strong and reliable effects on high frequency items. More immediately relevant to the current study is an experiment by Lee et al. (2005, Experiment 3) in which a reliable consistency effect was observed for high-frequency words.

Thus, the current findings leave some open questions regarding the strong frequency by consistency interaction and the null consistency effect for high frequency items. One possibility is that these effects arise from the types of confounds identified by Jared (1997). A later study (Jared, 2002) examined complex interactions among frequency, regularity and neighborhood properties and found results consistent with a significant frequency by consistency interaction with similar controls, suggesting against such an explanation. Another possibility is that the frequency manipulation in Study 1 was larger than those employed in studies reporting consistency effects for high frequency items. Consistency effects have been found most reliably for items with frequencies closer to 100 per million (Jared, 1997, 2002; Lee et al., 2005), yet in the current study, high frequency items appeared on average about 500 times per million words in the relevant corpora.

To extend and confirm the results of Study 1, it is critical to establish whether the model correctly predicts the interaction (or lack thereof) between frequency and consistency when stricter controls on neighborhood structure and a weaker manipulation of frequency are employed.

Study 2

It is an open empirical question whether the interaction between frequency and consistency in Chinese is an artifact of the same stimulus properties that appear to enhance this interaction in English. As a first step in exploring whether there are potential confounds between frequency and consistency in the lexicon at large, we conducted a corpus analysis. Each item in the training corpus (4468 items) from Modern Chinese Frequency Dictionary (1986) was tagged with

respect to the following dimensions: whether it is a phonogram, which phonetic component it contains, the pronunciation of the phonetic components, and the phonological relationship between phonetic component and character (in terms of regularity and consistency). Phonograms were identified based on the *XianDaiHanZiXingShengZiZiHui* (Ni, 1982). The number of characters and phonograms were counted as well as three types of phonograms: regular-consistent (R-C), regular-inconsistent (R-I) and irregular-inconsistent (I-I). As in Experiment 1, consistency level was calculated as the ratio of the summed frequency of friends to the summed frequency of the characters in a family, with family defined as all characters sharing the same phonetic component.

As summarized in Table 5, this corpus analysis revealed that consistency level is related to frequency, such that higher frequency items in general tend to be more consistent than low-frequency items. Taken together with the fact that there are many fewer high-frequency items than low-frequency items, and that a smaller proportion of these are phonograms than is the case for the low and middle frequency bands, it does seem plausible to suggest that the interaction between consistency and frequency is due in part to a confound with some additional statistical properties of the input corpus.

The comparisons of the consistency effect for high and low-frequency items may be confounded with differences in consistency level. Although Study 1 controlled for this possibility, as shown in Table 2, it did not control for other properties, such as family size or family frequency, nor for frequency of friends/enemies across levels of frequency for the consistency manipulation. Study 2 directly addresses these shortcomings by contrasting middle and low frequency items, thereby avoiding several of the confounds present in the high frequency stimuli, and enabling an examination of the interaction between consistency and frequency in more detail.

Connectionist model

Method—A set of 120 regular characters were selected from the training corpus, with 30 characters in each cell created by crossing frequency (middle/low) with consistency (consistent/inconsistent). Details of the stimulus properties are shown in Table 6. The simpler design permitted better control over family size across levels of consistency, in addition to a number of other stimulus properties such as family frequency, phonetic component frequency, and summed frequency of friends -- which were matched across levels of frequency. Finally, frequency levels were selected to be similar to Lee's (2005) and Jared's (1997) experiments 1&2 which showed consistency effect for high/medial frequency items and the weak or null interactions between frequency and consistency. Frequency counts were confirmed by drawing estimates from two corpora: *Modern Chinese Frequency Dictionary* (1986) and *Balanced Corpus of Modern Chinese* (Sun, 2006).

All 120 characters were tested on 40 models from Simulation 1. The testing method was the same as in the previous simulation. Both the SSE and naming accuracy were computed at 3 million training trials.

Result—Because of the lower overall frequency, the average SSE was nominally higher than simulation 1 and overall accuracy was lower (91% vs. 95% in simulation 1). The 2 (middle vs. low frequency) $\times$ 2 (consistent vs. inconsistent) ANOVA was conducted both for SSE and naming accuracy (Table 7). The overall correlation between SSE and response latency for all items in the behavioral experiment was 0.320 ($p < 0.01$).

A main effect of frequency was found both for SSE and naming accuracy. Middle frequency characters had lower SSE and fewer errors than low frequency characters. Consistent characters

were read more easily than inconsistent characters. Finally, a significant interaction between frequency and consistency was found for both SSE and accuracy.

For both accuracy and SSE, the simple consistency effect for middle frequency characters was significant in SSE by participants' analysis (treating each run of the model as a participant), but not in the item analysis, whereas the consistency effect was strongly significant for low-frequency items in both analyses. This replicates Lee's (2005) Experiment 3 results. The simple frequency effects for consistent and inconsistent characters were significant both for SSE and naming accuracy ($P_s < 0.01$). The frequency effect was larger for inconsistent characters ($SSE=0.59$) than for consistent characters ($SSE=0.29$).

Behavioral experiment

Method—Forty participants (none had participated in Experiment 1) were recruited from Beijing Normal University to name the new set of 120 characters used in Simulation 2. The procedure was the same as in Experiment 1.

Result—All participants' data were included and the overall naming accuracy was 92.19%. The reaction time for error responses was replaced with the participant's condition mean for correct trials. The average reaction time is presented in Figure 4. The ANOVA analysis for naming latency and accuracy is shown in Table 8.

The ANOVA results revealed that participants named middle-frequency characters significantly faster and accurately than low-frequency characters. Naming latencies were significantly longer for inconsistent characters than for consistent characters, and participants were significantly more accurate in naming consistent characters than inconsistent characters. As in Experiment 1, we found a significant interaction between frequency and consistency, both in the latency and accuracy data. The simple frequency effects both for consistent and inconsistent characters were significant ($P_s < 0.01$); the frequency effect was larger for inconsistent (124ms) than consistent characters (66ms). The consistency effect was 27ms for middle frequency characters and 84ms for low frequency characters. As predicted by Simulation 2, both for reaction time and naming accuracy, medium-frequency characters produced a significant consistency effect in the analysis by participants, but not by items. The consistency effect in low frequency characters was reliable both for participant and item analyses.

Analysis of Training Corpus and Hidden Units in the Model

In the following analysis, we examine the training corpus and the internal workings of the model to shed light on *how* these phenomena might emerge from the application of statistical learning mechanisms to the problem of mapping from print to sound in different writing systems.

Corpus Analysis—To examine the statistical properties that might give rise to functional units in reading Chinese characters, descriptive statistics for phonetic components and radicals were calculated in the training corpus: Pronunciations for all characters containing each radical and phonetic component were identified and counted in order to determine the reliability of these two sub-lexical structures in print to sound mappings.

For the 4468 characters in the training corpus, 298 constituent radicals and 1021 phonetic components were identified. Some radicals are also phonetic components, which creates some ambiguity with respect to what the functional units of the print-to-sound mapping might be. For example, as shown in Figure 1B, the radicals “阝” and “干” can function as phonetic components on their own, they also can be combined to form the phonetic component “皇”.

Most phonetic components are composed of more than two radicals, on average 2.4 radicals, with a range of one to six (Figure 5A). Figure 5B shows a histogram of the number of characters and syllables (including and not including tone) that can be formed by radicals and phonetic components. Radicals appear, on average, in about 39 characters with 27 different pronunciations each. In contrast, more than 70% of phonetic components are mapped to no more than four characters (on average about three), with an average of two distinct pronunciations. Thus, phonetic components, if they can be identified, are much stronger cues to pronunciation than radicals.

Emergent functional units in Chinese print-to-sound—In our model, no explicit coding was included in the representation to identify phonetic components, rather, the influence of these units emerged in the course of learning the mapping from orthographic to phonological representations. Our model did, however, replicate the empirical finding of a regularity/consistency effect, which is defined in terms of phonetic components. In order to explore the emergence of functional units for print to sound translation in the model, we analyzed the similarity space in hidden layer activations over the course of learning for a subset of items from the training set.

These analyses were carried out on a small subset of the training items, selected in order to probe whether the similarity space defined by hidden unit activations showed evidence of being shaped by phonetic components. A set of items that shared a phonetic component was selected, along with controls for orthographic and phonological similarity. In all, forty-four items were selected, of which 7 shared a phonetic component. In addition to the simple character, “台” (/t^haɪ/, three regular (台, 抬, 胎) and three irregular (治, 始, 始) items from the same family were included. As a comparison group, a set of 27 orthographically similar items was selected to have the same degree of orthographic similarity to the critical items as they would if they shared a phonetic component. This was done by finding items that shared two radicals that did not comprise a phonetic. For example, 著 shares the radicals 艹 and 者 with the target item 著 (comprising 艹, 者 and 者). In addition, we selected phonologically similar characters, which were 10 high frequency homophones of two irregular characters (治, 始), included in order to address the role of phonology by itself in organizing the hidden layer.

Each character was clamped on the input layer for 12 time ticks and the activation values of 200 hidden units were computed at the last time tick. Multidimensional scaling was computed to characterize the similarity space both on the orthographic and hidden layers. Tests for similarity of groups of items (shared phonetic component, orthographic controls, and phonological controls) were conducted on the Euclidean distances within and between groups in this similarity space.

Results from the Euclidean distance measures based on the similarity space of characters derived from hidden unit activations are plotted over the course of training in Figure 6A. At the beginning of training, overall activity levels on the hidden units were quite small (because of the small, random weights to which the model was initialized), resulting in relatively small distances driven entirely by the input. As the model learned to map from print to sound, however, differences emerged, such that by the end of training, items that share a phonetic component were on average more similar to one another than they were to either orthographic or phonological controls. This was supported by a 2 (training period: initial weights, after 3 million trials) by 4 (stimulus comparison: within phonetic component's family, orthographic control, phonological control, between phonetic component's family) ANOVA, which demonstrated a main effect of training, $F(1,313) = 553.68$, $MSE = 50.58$, $p < 0.01$; a main effect of stimulus, $F(3,313) = 22.17$, $MSE = 3.95$, $p < 0.01$, and a significant training by stimulus interaction, $F(3,313) = 10.78$, $MSE = 0.99$, $p < 0.01$. Planned comparisons at 3 million training trials revealed that within-phonetic component's family comparisons had smaller

distances (1.41) than the other three conditions: Orthographic controls (1.90, $t(46) = 3.1$, $p < 0.01$), phonological controls (2.03, $t(29) = 2.49$, $p < 0.01$) or between-family comparisons (2.13, $t(278) = 6.58$, $p < 0.01$).

For comparison to the MDS solutions shown for the hidden units, Panel B of Figure 6 shows a solution for the similarity space represented on the orthographic input. Note that, although the members of the same family form a loose cluster, for many items, there are non-family members that are much closer in similarity space than other family members. In Panel C, it is clear that in the initially randomized state, the hidden unit activations essentially recapitulate the input representations. Indeed, correlations between pairwise distances on the input and hidden units before training are nearly perfect ($r = 0.99$, $p < 0.001$). After 3 million trials of training (Panel D) items with the same phonetic component (ㄜ/ $t^h a I$ /) form a relatively tight cluster. Within this cluster, there are also sub-clusters of items that share a pronunciation, so that, for example, the items with the same pronunciation (/ㄜ/ $t^h a I$ /: 拈, 砵, 腩) and the items with different pronunciations (拈/ i /, 治/ $t s a I$ /, 始/ $s a I$ /) occupy distinct regions of the similarity space. The correlation between input and hidden unit representations ($r = 0.48$, $p < 0.001$) is smaller than what was observed initially ($t(943) = 54.35$, $p < 0.001$).

The relatively high degree of similarity among these items is not simply a reflection of their output similarity. A number of orthographically related homophones (拈/ i /: 仅移姚迺宜; 治/ $t s a I$ /: 致志置制质) are quite distant from this cluster in the similarity space. The distance between homophones (mean distance = 2.03, STD=0.52, N=10) and irregular items (/ㄜ/ $t^h a I$ /) was significantly larger ($t(20) = 3.11$, $p < 0.01$) than the distance to other phonetic component's family members (mean distance = 1.37, STD=0.48, N=12).

Figure 7 shows hidden unit activations for a set of items from the similarity space analyses. The overlap among items that share the phonetic component (ㄜ) and are pronounced /ㄜ/ $t^h a I$ / is strikingly compact. Two hidden units are near their maximum activation for all three of these items. One of these same hidden units is active for all items that contain the phonetic component (ㄜ), independent of how they are pronounced. In contrast, among items that are equivalently similar in terms of their orthographic representations, but do not share a phonetic component, there is no such overlap in the hidden layer representation. Thus, we can see that functional sub-lexical units are represented in a relatively compact manner in the model.

It is interesting to contrast the patterns of hidden unit representation to what has been observed in similar analyses of English reading models (Harm & Seidenberg, 1999; Harm, McCandliss & Seidenberg, 2003). In those models, individual words evoke much more diffuse and widespread activity over the hidden layer, likely reflecting both the denser representations permitted by the much smaller range of possible inputs, and the contribution of smaller sub-lexical units (e.g., single letters mapping with some consistency onto single phonemes). In Chinese, the relatively large grain size and high degree of arbitrariness require a much sparser representation. This may also explain why a larger number of hidden units are needed for Chinese than for English in these models.

Items that share the same phonetic component are treated as more similar than items that share the same amount of orthographic information, even when these map on to highly distinct phonological outputs. For example, the characters 拈 and 治 have essentially *no* overlap in their pronunciations (/ㄜ/ $t^h a I$ /, /ㄜ/ $t s a I$ /), and yet they share similar representations (in this case, largely defined by a single hidden unit). This example clearly illustrates that the representations arrived at in the acquisition of print-to-sound mappings are organized by both the information-bearing structures in the input and the similarity of their mapping to the output, reflecting the model's extraction of sub-lexical regularities in the translation from print to sound. The model thus demonstrates how compact and superpositional representations of shared orthographic

structure may be learned and play a role in reading aloud, even when the grain size of mappings from print to sound is fairly coarse, and the mappings themselves are probabilistic with a high degree of arbitrariness.

Discussion

This study explored the interactions between consistency and frequency under tight stimulus control and manipulation of “medium” vs. “low” frequency, in order to determine the robustness of this finding. As in previous studies with similar manipulations, the interaction between frequency and consistency was much weaker than in Study 1 (but nonetheless reliable), and a consistency effect was found for the medium-frequency items (although this was only significant by subjects). Thus, the interaction between frequency and consistency in Study 1 was not the result of stimulus factors that could not be controlled due to the complexity of the design. The predictions of the model were confirmed by behavioral experiments on the same testing items; further analyses revealed that the functional units on which the frequency by consistency interaction depends are emergent properties of the application of the statistical learning rules of the model to the corpus.

Analyses of the hidden units revealed that these effects were driven by the organization of representations that encode probabilistic information about the functional units that support mappings from orthography to phonology. Further, they demonstrate how these representations might emerge over development as a result of statistical learning, and without assumptions about the level at which regularities exist in the writing system. Whereas at first, the similarity space defined by hidden unit activations is organized by orthographic similarity, relatively early in training, words that share a phonetic component begin to be represented as more similar to one another than orthographic or phonological controls.

There is some empirical evidence to support the model’s account of learning to read in Chinese as well: Shen and Bear (2000) collected invented spellings from schoolchildren and classified them according to whether they contained orthographic, phonological or semantic errors. They found that the earliest invented spellings involved mostly orthographic confusions, whereas by fourth grade, the majority of invented spellings involved insertion of a phonetic radical that was consistent with the intended character’s pronunciation.

The emergence of regularity and consistency effects in children is also broadly consistent with this analysis. Early readers (second graders) show a strong regularity effect, tending to misread phonograms according to the pronunciation of their phonetic component when it appears by itself (Shu, Anderson, & Wu, 2000). By sixth grade, however, subtler effects of phonetic family (i.e., consistency effects) are also found (Chen, Shu, Wu, & Anderson, 2003; Yang, & Peng, 1997). Analyses of children’s reading materials (Shu, Chen, Wu, & Xuan, 2003) suggest that this transition is driven in part by changes in the statistics of children’s reading materials. Phonetic components are learned as individual words in the early grades, whereas in later grades, families of words that share a phonetic component are learned as the written vocabulary expands to include a larger number of phonograms. A more detailed simulation of this developmental pattern could be achieved with incremental training (Powell, Plaut, & Funnell, 2006).

General Discussion

A computational model of reading aloud in Chinese based on connectionist principles correctly simulated effects of regularity, consistency, frequency and their interaction, as observed in a behavioral study of reading aloud. In both the human and simulation data, regularity and consistency effects were each larger for low- than for high-frequency words. A regularity effect was found for inconsistent items, a finding that would seem on its face inconsistent with this

modeling framework. The simulation correctly captured this effect, however, demonstrating that it can be explained without rule-based processing or the inclusion of phonetic components as a special “level” of the representation. A more in-depth exploration of the frequency by consistency interaction demonstrated that it was robust to strict controls of neighborhood characteristics, and could be found even under a relatively weak manipulation of frequency. The model correctly simulated these findings as well. Finally, analyses of the training set and internal representations in the model revealed that its sensitivity to consistency and regularity results from the emergence of phonetic components as functional units over the course of training.

Simulating language-general and language-specific effects in reading Chinese

The current simulations applied the same basic architecture and learning rules used to simulate a variety of phenomena in English reading (Harm & Seidenberg, 1999, 2004; Seidenberg & Zevin, 2006), to capture both effects that are common to both English and Chinese – i.e., frequency, consistency and their interaction – and effects that are specific to a particular writing system – i.e., regularity as defined with reference to phonetic components in Chinese. In both cases, the explanation rests on the operation of domain-general statistical learning mechanisms.

Frequency and consistency effects are explained by complementary features of the learning algorithm applied in the current simulation. On the one hand, learning in the models is local, so that practice with a particular item has the greatest impact on performance for that item itself. On the other hand, because representation in the model is superpositional – i.e. many different items are represented as patterns of activation over the same units – the model has an inherent tendency to map similar inputs onto similar outputs. Thus, whereas frequency effects reflect previous experience with a particular word, consistency effects reflect experience with all similar words. The effects of these properties interact because sufficient practice with a particular word can override ambiguity arising from similarly spelled words. These same principles hold, even as the definition of consistency itself (Jared, 2002; Treiman et al., 2006) is further refined, because the critical assumption is the role of statistics at multiple grain sizes, not the particular unit (letter, grapheme, word body or whole word) over which the statistics are relevant.

Studies of the interaction between consistency and frequency – including those presented here -- typically take advantage of the power of factorial designs. Nonetheless, it is clear that both consistency and frequency are essentially continuous variables. Thus it is fair to suggest that the strength of the interaction between them will depend on the details of how each is manipulated. Indeed, some of the contradictory findings across studies reviewed by Jared (1997, 2002) are apparently the result of quantitative differences in the neighborhood structure of inconsistent words. Words like TINT, that have many similarly pronounced, relatively high-frequency “friends” (HINT, MINT, LINT, etc.) and few, relatively low-frequency “enemies” (PINT) show relatively weak or null consistency effects, even when they are relatively low frequency.

Less attention has been paid to the role of the strength of the frequency manipulation in modulating this interaction. The studies reviewed and presented here suggest that the interaction between frequency and consistency is strongest for relatively large frequency manipulations. This is correctly predicted by the model because, as a statistical learning model, it naturally takes into account the continuous nature of both frequency and consistency. Thus the model’s prediction of a frequency by consistency interaction for a particular set of items is not the same as a blanket prediction from a verbal model that these variables should interact under all conditions. That the model fits both strong and weak interactions of frequency and consistency revealed across human studies (and could conceivably fit null effects with the

appropriate stimulus manipulations) suggests that humans are sensitive to the same statistical properties that drive its performance.

Regularity effects based on phonetic component pronunciation are unique to the Chinese writing system, and yet they can be explained as resulting from the same properties of statistical learning. Most phonetic components can appear on their own as individual characters (A few bound phonetics are archaic characters whose pronunciation are lost in modern Chinese -- see Shu et al., 2003 -- however all phonetic components of the characters used in the current study were unbound). This creates a very effective learning situation for mapping from a phonetic component to its canonical pronunciation. In terms of the learning algorithm used here, all of the error on learning trials when a phonetic component is presented in isolation accrues to strengthening the connections between that component and its pronunciation. In this way, the canonical pronunciation of each phonetic component attains something of a “special” status, such that characters whose pronunciations deviate from it are more difficult to read. Critically, the special status of phonetic components in these models is driven by statistics of the input rather than by an *a priori* representational scheme imposed by the modelers.

The emergence of functional units and “sub-lexical” processing in print-to-sound conversion for Chinese

Previous models of Chinese reading have included separate representational “layers” for radicals and characters (Perfetti, Liu, & Tan, 2005) or designed input representations explicitly around the radical as an unanalyzed, holistic unit (Hsiao & Shillcock, 2005, 2006) and in one case (Xing, Shu, & Li, 2002, 2004) included phonetic components along with their pronunciations as part of the input to the model in order to capture regularity and consistency effects (but see Chen & Peng, 1994). In contrast, the current model uses a distributed representation of Chinese characters that eschews any *a priori* assumptions about the functional units that underlie spelling-to-sound translation.

In some sense, the current coding scheme for the input layer contains some language-specific features -- a more elaborate model that learned to identify contrastive features from an input that incorporated features of the visual system would be truly language-general, in the sense that it could take either English or Chinese as its input. Nevertheless, the current model is language-general in the sense that its orthographic representation is based on the smallest contrastive unit in the writing system (strokes), and organized to capture the visual similarity of characters, rather than fixing a specific level of representation such as the radical or phonetic component as a functional unit. This is illustrated in Figure 6B, which shows that the similarity space on the orthographic input layer is loosely organized by orthographic units at multiple levels of representation. For example, the cluster of characters in the bottom left hand corner mostly share the 氵 (“water”) radical, but some also contain a 心 (“heart”) radical (which also has three strokes and appears on the left). In response to the demands of learning to map from print-to-sound, the model’s representations are increasingly organized by behaviorally relevant structures, i.e., phonetic components (Figures 6C and 6D). Thus, it is not necessary to assume a fixed level of orthographic representation in order to simulate print-to-sound conversion in Chinese.

Previous studies have used priming paradigms to address the issue of whether radicals have special status in the organization of the adult reading system (Ding, Peng, & Taft, 2004; Taft & Zhu, 1997; Taft, Zhu, & Ding, 2000; Taft, Zhu, & Peng, 1999). The results of these studies have shown that radicals influence lexical access, and may be a basic unit of Chinese orthographic representation (Ding, Peng, & Taft, 2004). The analyses presented in Study 2, however, suggest that radicals are an unlikely candidate for a functional unit in mapping from spelling to sound. It may be that they are more consistently related to semantics, or that they are important in some early stage of character identification.

Although it is beyond the scope of the current work to simulate the priming effects that are taken as evidence for radicals as functional units, the model provides an alternative perspective on how functional units emerge over the development of reading, instead of assuming a fixed functional unit for all reading processing. A topic for future research is the emergence of similar structure in mappings from orthography to semantics. Just as the phonetic component of a phonogram provides cues to its pronunciation, the semantic component provides cues to meaning. For example, the radical for water (氵) appears in many phonograms that are related to water (湖, *lake*, 河, *river*, 渴, *thirsty*, 游, *swim*), but it can also appear in phonograms that are not (法, *law*, 洽, *negotiate*). This radical may thus be viewed as encoding probabilistic information about semantics at the sub-lexical level.

The role of phonology in Chinese reading

A number of studies have focused on the role of phonology in computing meaning from sound (e.g., Perfetti & Tan, 1998; Perfetti & Zhang, 1995). This work led to the Lexical Constituency Model (LCM, Perfetti, Liu, & Tan, 2005), which simulates the time-course of effects of graphical, phonological and semantic primes. The central theoretical point of the LCM is that phonology is rapidly and obligatorily computed during reading (even reading for meaning). While there are some data inconsistent with this claim (Zhou & Marslen-Wilson, 2000), there are a large number of phenomena that appear difficult to explain in any other way (Chen, Flores d'Arcais, & Cheung, 1995; Chen & Shu, 2001; Xu, Pollatsek, & Potter, 1999; Zhou & Marslen-Wilson, 1999b).

Because we present a model of the print-to-sound, we cannot directly address the role of phonology in the computation of semantics. A model of English reading with a similar architecture for print-to-sound that also included semantics (Harm & Seidenberg, 2004), was able to correctly simulate a wide range of phonological effects in the computing semantics from print. Given the gross differences in the ease of print-to-sound mapping, and the probabilistic cues to meaning that distinguish Chinese from alphabetic writing systems, we expect the division of labor between direct and phonologically-mediated mappings from print to semantics to differ between the two languages. The current model, however, demonstrates that print to sound mappings can be learned and processed efficiently independently of semantics, and thus could plausibly contribute to semantic activation in a model that included semantics.

In contrast to the LCM, however, we have specified a *sub-lexical* mechanism for print-to-sound conversion. In the LCM, it is assumed that each character has a stored pronunciation, and these are simulated in practice by hand-coding the pronunciations into the model. The current model learns print-to-sound mappings over distributed representations, and thus is capable of discovering regularities among sub-lexical units of varying size and consistency. The behavioral results presented here (and Hsu, Tsai, Lee, & Tzeng, 2009; Fang, Horng, & Tzeng, 1986; Hue, 1992; Lee, Tsai, Su, Tzeng, & Hung, 2005; Peng & Yang, 1997) demonstrate that these sub-lexical regularities play an important role in a task that directly taps print-to-sound conversion, and their simulation in the model demonstrates that they can be explained in terms of sub-lexical processing.

Conclusion

We have adapted a statistical learning model of reading aloud initially developed to study an alphabetic writing system (English) to an ideographic one (Chinese). This was possible because the central assumptions of the model are simple and readily generalize across writing systems. The success of this adaptation suggests that the same framework can be applied to understanding reading skill across a wide variety of languages, despite gross differences in their surface properties. This is a critical first step in providing a mechanistic explanation of

many cross-linguistic phenomena, including the differential impact of factors that predict reading success (McBride-Chang, Chao, Liu, et al., 2005; Shu, Peng, & McBride-Chang, 2008), differences in the prevalence of developmental disorders of reading (Johansson, 2006; Shu, McBride-Chang, Wu, & Liu, 2006; Shu, Meng, Chen, Luan, & Cao, 2005) and patterns of reading disorder subsequent to brain injury (Bi, Han, Weekes, & Shu, 2007; Jefferies, Sage, & Ralph, 2007; Woollams, Lambon Ralph, Plaut, & Patterson, 2007).

The central explanatory principles of this framework may be usefully expanded in future research to examine the unique aspects of semantic encoding in Chinese. Similar types of probabilistic mappings are present in print to *meaning* in Chinese, and future research should focus on incorporating this very unusual property of the writing system. Preliminary studies suggests that it may be possible to capture the differential division of labor between semantically mediated and direct translation from print to sound, and the differential contributions of phonological and semantic processing across languages in this same framework (Yang, McCandliss, Shu, & Zevin, 2008; Yang, Zevin, Shu, McCandliss, & Li, 2006). Thus, this framework has the potential to explain how the same functional architecture can give rise to both strikingly similar (e.g., the case of regularity and consistency effects) and highly distinct (e.g., differential contributions of phonological abilities to reading acquisition) outcomes across writing systems.

References

- Balota D, Ferraro F. A dissociation of frequency and regularity effects in pronunciation performance across young adults, older adults, and individuals with senile dementia of the Alzheimer type. *Journal of Memory and Language* 1993;32(5):573–592.
- Bentin S, Frost R. Processing lexical ambiguity and visual word recognition in a deep orthography. *Memory & Cognition* 1987;15(1):13–23.
- Bi Y, Han Z, Weekes B, Shu H. The interaction between semantic and the nonsemantic systems in reading: Evidence from Chinese. *Neuropsychologia* 2007;45(12):2660–2673. [PubMed: 17568630]
- Chen H-C, Flores d'Arcais GB, Cheung S-L. Orthographic and phonological activation in recognizing Chinese characters. *Psychological Research* 1995;58(2):144–153.
- Chen H-C, Shu H. Lexical activation during the recognition of Chinese characters: Evidence against early phonological activation. *Psychonomic Bulletin & Review* 2001;8(3):511–518. [PubMed: 11700902]
- Chen X, Shu H, Wu N, Anderson RC. Stages in learning to pronounce Chinese characters. *Psychology in the Schools* 2003;40(1):115–124.
- Chen, Y.; Peng, D. A connectionist model of recognition and naming of Chinese characters. In: Chang, H-W.; Huang, J-T.; Hue, C-W.; Tzeng, O., editors. *Advances in the study of Chinese language processing*. Vol. 1. Taipei: National Taiwan University Press; 1994. p. 211-240.
- Chinese National Language Affairs commission. *Chinese Character Component Standard of GB13000.1 Character Set for information Processing*. 1997
- Coltheart M, Curtis B, Atkins P, Haller M. Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review* 1993;100(4):589–608.
- 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(1):204–256. [PubMed: 11212628]
- DeFrancis, J. *Visible Speech: The Diverse Oneness of Writing Systems*. Honolulu, HI: University of Hawaii Press; 1989.
- Ding G, Peng D, Taft M. The Nature of the Mental Representation of Radicals in Chinese: A Priming Study. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2004;30(2):530–539.
- Fang, SP.; Horng, RY.; Tzeng, OJL. Consistency effects in the Chinese character and pseudo-character naming tasks. In: Kao, HSR.; Hoosain, R., editors. *Linguistics, psychology, and the Chinese language*. University of Hong Kong: Center of Asian Studies; 1986. p. 11-21.
- Feldman LB, Siok WW. Semantic Radicals Contribute to the Visual Identification of Chinese Characters. *Journal of Memory and Language* 1999;40(4):559–576.

- Forster KI, Forster JC. DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, & Computers* 2003;35:116–124.
- Frost R, Katz L, Bentin S. Strategies for visual word recognition and orthographical depth: a multilingual comparison. *Journal of Experimental Psychology: Human Perception and Performance* 1987;13(1): 104–115. [PubMed: 2951484]
- 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(3):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(3):662–720. [PubMed: 15250780]
- Havelka J, Rastle K. The assembly of phonology from print is serial and subject to strategic control: Evidence from Serbian. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2005;31(1):148–158.
- Hinton GE, Shallice T. Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review* 1991;98(1):74–95. [PubMed: 2006233]
- Hsiao JH, Shillcock R. Foveal splitting causes differential processing of Chinese orthography in the male and female brain. *Cognitive Brain Research* 2005;25(2):531–536. [PubMed: 16154326]
- Hsiao, JH-w; Shillcock, R. Hemispheric Differences Emerge from Perceptual Learning: Evidence from Modeling Chinese Character Pronunciation; Paper presented at the Proceedings of the Twenty eighth Annual Conference of the Cognitive Science Society; Vancouver, BC, Canada. 2006.
- Hsu C-H, Tsai J-L, Tzeng OJ. Orthographic combinability and phonological consistency effects in reading Chinese phonograms: An event-related potential study. *Brain and Language* 2009;108(1): 56–66. [PubMed: 18951624]
- Huang, B.; Liao, X. *XianDaiHanYu*(3rd supplement). Beijing: Higher Education Press; 2002.
- Hue, C. Recognition processes in character naming. In: Chen, H.; Tzeng, OJL., editors. *Language Processing in Chinese*. North-Holland: Elsevier Science Publisher; 1992. p. 93-107.
- Jared D. Spelling-Sound Consistency Affects the Naming of High-Frequency Words. *Journal of Memory and Language* 1997;36(4):505–529.
- Jared D. Spelling-sound consistency and regularity effects in word naming. *Journal of Memory and Language* 2002;46(4):723–750.
- Jared D, McRae K, Seidenberg M. The basis of consistency effects in word naming. *Journal of Memory and Language* 1990;29:687–715.
- Jefferies E, Sage K, Ralph MAL. Do deep dyslexia, dysphasia and dysgraphia share a common phonological impairment? *Neuropsychologia* 2007;45(7):1553–1570. [PubMed: 17227679]
- Johansson B. Cultural and linguistic influence on brain organization for language and possible consequences for dyslexia: A review. *Annals of Dyslexia* 2006;56(1):13–50. [PubMed: 17849207]
- Kang H, Simpson GB. Local strategic control of information in visual word recognition. *Memory & Cognition* 2001;29:648–655.
- Lee C-Y, Tsai J-L, Su EC-I, Tzeng OJL, Hung DL. Consistency, regularity, and frequency effects in naming Chinese characters. *Language and Linguistics* 2005;6(1):75–107.
- Lee, I. *The Korean Language* (transl. Robert Ramsey). Albany, NJ: State University of New York Press; 2000.
- Leong, CK. Paradigmatic analysis of Chinese word reading: Research findings and classroom practices. In: Leong, CK.; Joshi, RM., editors. *Cross-language studies of learning to reading and spell: Phonological and orthographic processing*. Dordrecht: Kluwer Academic Publishers; 1997. p. 379-417.
- Li, Y.; Kang, J. *Information analysis of usage of characters in Modern Chinese*. Shanghai: Shanghai Education Publisher; 1993. Analysis of phonetics of the ideophonetic characters in Modern Chinese; p. 84-98.
- Malone K. Benjamin Franklin on spelling reform. *American Speech* 1925;1(2):96–100.

- Mattingly II. Morphological structure and segmental awareness. *C.P.C.: Cahiers de Psychologie Cognitive European Bulletin of Cognitive Psychology* 1987;7:488–493.
- McBride-Chang C, Cho J-R, Liu H, Wagner RK, Shu H, Zhou A, et al. Changing models across cultures: Associations of phonological awareness and morphological structure awareness with vocabulary and word recognition in second graders from Beijing, Hong Kong, Korea, and the United States. *Journal of Experimental Child Psychology* 2005;92(2):140–160. [PubMed: 15904930]
- McClelland JL, Rumelhart DE. An interactive activation model of context effects in letter perception: Part I. An account of Basic Findings. *Psychological Review* 1981;88:375–407.
- Modern Chinese Frequency Dictionary. Beijing: Beijing Language Institute Publisher; 1986. (in Chinese)
- Ni, H. *XianDaiHanZiXingShengZiZiHu*. Beijing: Language and Culture Press; 1982.
- Pearlmutter BA. Gradient calculations for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks* 1995;6(5):1212–1228. [PubMed: 18263409]
- Peng D, Yang H. The phonological processing of Chinese phonograms. *Asia and Pacific Journal of Speech, Language and Hearing* 1997;2:177–195.
- Perfetti CA, Liu Y, Tan LH. The Lexical Constituency Model: Some implications of research on Chinese for general theories of reading. *Psychological Review* 2005;112(1):43–59. [PubMed: 15631587]
- Perfetti CA, Tan LH. The Time Course of Graphic, Phonological, and Semantic Activation in Chinese Character Identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1998;24(1):101–118.
- Perfetti CA, Zhang S. Very Early Phonological Activation in Chinese Reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1995;21(1):24–33.
- Plaut, DC. Connectionist approaches to reading. In: Snowling, M.; Hulme, C.; Seidenberg, MS., editors. *The science of reading: A handbook*. Blackwell: Oxford; 2005. p. 24–38.
- Plaut DC, McClelland JL, Seidenberg MS, Patterson K. Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review* 1996;103(1):56–115. [PubMed: 8650300]
- Plaut DC, Shallice T. Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology* 1993;10(5):377–500.
- Powell D, Plaut DC, Funnell E. Does the PMSP connectionist model of single word reading learn to read in the same way as a child? *Journal of Research in Reading* 2006;29(2):229–250.
- Raman I, Baluch B, Besner D. On the control of visual word recognition: Changing routes versus changing deadlines. *Memory & Cognition* 2004;32:489–500.
- Rayner K, Foorman BR, Perfetti CA, Pesetsky D, Seidenberg MS. How Psychological Science Informs the Teaching of Reading. *Psychological Science in the Public Interest* 2001;2:31–74.
- Seidenberg MS. The time course of phonological code activation in two writing systems. *Cognition* 1985;91(1):1–30. [PubMed: 4039640]
- Seidenberg, MS. Beyond orthographic depth: Equitable division of labor. In: Frost, BR.; Katz, LL., editors. *Orthography, phonology, morphology, and meaning*. North-Holland: Springer-Verlag; 1992. p. 85–118.
- Seidenberg MS, McClelland JL. A distributed, developmental model of word recognition and naming. *Psychological Review* 1989;96(4):523–568. [PubMed: 2798649]
- Seidenberg, MS.; Plaut, DC. Progress in understanding word reading: Data fitting versus theory building. In: Andrews, S., editor. *From inkmarks to ideas: Current issues in lexical processing*. Hove, UK: Psychology Press; 2006. p. 25–49.
- Seidenberg, MS.; Zevin, JD. Connectionist models in developmental cognitive neuroscience: Insights about critical periods. In: Munakata, Y.; Johnson, MH., editors. *Attention and Performance XXI: Processes of Change in Brain and Cognitive Development*. Oxford University Press; 2006. p. 585–612.
- Share DL. On the Anglocentricities of Current Reading Research and Practice: The Perils of Overreliance on an “Outlier” Orthography. *Psychological Bulletin* 2008;134(4):584–615. [PubMed: 18605821]
- Shen HH, Bear DR. Development of orthographic skills in Chinese children. *Reading and Writing* 2000;13(3):197–236.

- Shu H, Anderson RC, Wu N. Phonetic awareness: Knowledge of orthography-phonology relationships in the character acquisition of Chinese children. *Journal of Educational Psychology* 2000;92(1):56–62.
- Shu H, Chen X, Anderson RC, Wu N, Xuan Y. Properties of school Chinese: Implications for learning to read. *Child Development* 2003;74(1):27–47. [PubMed: 12625434]
- Shu H, McBride-Chang C, Wu S, Liu H. Understanding Chinese developmental dyslexia: Morphological awareness as a core cognitive construct. *Journal of Educational Psychology* 2006;98(1):122–133.
- Shu H, Meng X, Chen X, Luan H, Cao F. The subtypes of developmental dyslexia in Chinese: evidence from three cases. *Dyslexia* 2005;11(4):311–329. [PubMed: 16355750]
- Shu H, Peng H, McBride-Chang C. Phonological awareness in young Chinese children. *Developmental Science* 2008;11(1):171–181. [PubMed: 18171377]
- Shu, H.; Wu, N. Growth of orthography-phonology knowledge in the Chinese writing system. In: Li, P.; Bates, E.; Tan, L.H.; Tzeng, O.J.L., editors. *Handbook of East Asian Psycholinguistics*. Vol. 1. Cambridge: Cambridge University Press; 2006. p. 103–113. Chinese Psycholinguistics
- Shu H, Zhang H. The process of pronouncing the Chinese character by adult skilled reader. *Acta Psychologica Sinica* 1987;19(3):282–290. (in Chinese).
- Sun, M. Balanced corpus of modern Chinese. Tsinghua University: AI and NLP Group; 2006. (in Chinese)
- Taft, M. Processing of characters by native Chinese readers. In: Li, P.; Bates, E.; Tan, L.H.; Tzeng, O.J.L., editors. *Handbook of East Asian Psycholinguistics*. Vol. 1. Cambridge: Cambridge University Press; 2006. p. 237–249. Chinese Psycholinguistics
- Taft M, Zhu X. Submorphemic processing in reading Chinese. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1997;23(3):761–775.
- Taft M, Zhu X, Ding G. The relationship between character and radical representation in Chinese. *Acta Psychologica Sinica* 2000;32:1–12.
- Taft M, Zhu X, Peng D. Positional specificity of radicals in Chinese character recognition. *Journal of Memory and Language* 1999;40(4):498–519.
- Taraban R, McClelland J.L. Conspiracy effects in word pronunciation. *Journal of Memory and Language* 1987;26(6):608–631.
- Treiman R, Mullennix J, Bijeljac-Babic R, Richmond-Welty E. The special role of rimes in the description, use, and acquisition of English orthography. *Journal of Experimental Psychology: General* 1995;124(2):107–136. [PubMed: 7782735]
- Treiman R, Kessler B. Spelling as statistical learning: Using consonantal context to spell vowels. *Journal of Educational Psychology* 2006;98(3):642–652.
- Treiman R, Kessler B, Bick S. Context sensitivity in the spelling of English vowels. *Journal of Memory and Language* 2002;47:448–468.
- Treiman R, Kessler B, Zevin JD, Bick S, Davis M. Influence of consonantal context on the reading of vowels: Evidence from children. *Journal of Experimental Child Psychology* 2006;93(1):1–24. [PubMed: 16115645]
- Venezky, R.L. *The structure of English orthography*. The Hague: Mouton; 1970.
- Venezky, R.L. *American Way of Spelling: The Structure and Origins of American English Orthography*. New York: The Guildford Press; 1999.
- Woollams AM, Lambon Ralph MA, Plaut DC, Patterson K. SD-Squared: On the association between semantic dementia and surface dyslexia. *Psychological Review* 2007;114(2):316–339. [PubMed: 17500629]
- Xing, H.; Shu, H.; Li, P. A Self-Organizing connectionist model of character acquisition in Chinese. In: Gray, W.D.; Schunn, C.D., editors. *Proceedings of the Twenty Fourth Annual Conference of the Cognitive Science Society*; Mahwah, NJ: Lawrence Erlbaum; 2002. p. 950–955.
- Xing H, Shu H, Li P. The acquisition of Chinese characters: Corpus analyses and connectionist simulations. *Journal of Cognitive Science* 2004;5(1):1–49.
- Xu Y, Pollatsek A, Potter MC. The Activation of Phonology During Silent Chinese Word Reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1999;25(4):838–857.

- Yang, H.; Peng, D. The learning and naming of Chinese characters of elementary school children. In: Chen, HC., editor. *Cognitive Processing of Chinese and Related Asian Languages*. Hong Kong: The Chinese University Press; 1997. p. 323-346.
- Yang, J.; McCandliss, BD.; Shu, H.; Zevin, JD. Division of Labor between Semantics and Phonology in Normal and Disordered Reading Development across Languages; Paper presented at the Proceedings of the 30th Annual Conference of the Cognitive Science Society; Washington DC, USA. 2008.
- Yang, J.; Zevin, JD.; Shu, H.; McCandliss, BD.; Li, P. A "Triangle" Model of Chinese Reading; Paper presented at the Proceedings of the Twenty eighth Annual Conference of the Cognitive Science Society; Vancouver, Canada. 2006.
- Zevin JD, Seidenberg MS. Age of acquisition effects in word reading and other tasks. *Journal of Memory and Language* 2002;47(1):1-29.
- Zevin JD, Seidenberg MS. Simulating consistency effects and individual differences in nonword naming: A comparison of current models. *Journal of Memory and Language* 2006;54(2):145-160.
- Zhang S, Perfetti CA, Yang H. Whole Word, Frequency-General Phonology in Semantic Processing of Chinese Characters. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1999;25(4):858-875.
- Zhou X, Marslen-Wilson W. The nature of sublexical processing in reading Chinese characters. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1999a;25(4):819-837.
- Zhou X, Marslen-Wilson W. Phonology, Orthography, and Semantic Activation in Reading Chinese. *Journal of Memory and Language* 1999b;41(4):579-606.
- Zhou X, Marslen-Wilson W. The Relative Time Course of Semantic and Phonological Activation in Reading Chinese. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2000;26(5):1245-1265.
- Zhu, X. Analysis of cueing function of phonetic components in modern Chinese. In: Yuan, X., editor. *Proceedings of the symposium on the Chinese language and characters*. Beijing, China: Guang Ming Daily Press; 1988. p. 85-99.
- Ziegler JC, Goswami U. Reading acquisition, developmental dyslexia, and skilled reading across languages: A psycholinguistic grain size theory. *Psychological Bulletin* 2005;131(1):3-29. [PubMed: 15631549]
- Zorzi M, Houghton G, Butterworth B. Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance* 1998;24(4):1131-1161.

Acknowledgments

The authors would like to thank Ping Li and Hongbing Xing for contribution on the orthographic representation, Haiyan Zhou, Youyi Li and Xiaojuan Wang's work on the empirical data, and Mike Harm for technical assistance, the Mikenet software and interesting discussions. This research was supported by NSF of China 30870758, 60534080 and NSF of Beijing 7092051 (HS) grants, NSF REC 0337765 (BDM).

Figure 1.

A. Two orthographic “neighborhoods” in Chinese. Examples on the left are from a perfectly consistent neighborhood: Every word containing the phonetic component is pronounced the same way. On the right, an inconsistent neighborhood is shown. The top two items are “regular,” in that they share the pronunciation of the phonetic component when it appears as a simple character. The bottom three items have different pronunciations from the component as a simple character and are therefore irregular. B. Two radicals in phonetic component of Panel A (left) also can be phonetic components to form their families.

Figure 2. Architecture and representational scheme. A. The architecture of the model; arrows indicate connections between every unit in each layer. B. The representation of a Chinese syllable example on the phonological output layer. C. The representation of a Chinese character including 21 units for overall structure and 249 units to represent 7 radical slots.

Figure 3.

Predictions from the model for all three item types: Regular-consistent (R-C), Regular-inconsistent (R-I) and Irregular-inconsistent (I-I) at two levels of frequency (high and low) are shown in the graph on the left. On the right are response latency data from Chinese adult readers naming the same items. The model correctly predicts that both consistency and regularity interact with frequency, and a regularity effect for inconsistent items.

Figure 4. Predictions from the model for interaction between consistency (consistent and inconsistent) and frequency (middle and low) are shown in the graph on the left. On the right are response latency data from Chinese adult readers naming on the same characters. A strong consistency effect was found for low frequency and the weak consistency effect was found in high frequency for both participants and models.

Figure 5.

Panel A is a histogram showing how many phonetic components are composed of different numbers of radicals -- most are composed of more than two. Panels B, C, and D are histograms showing the number of characters (B) syllables (C) and syllables not counting tone (D) in which radicals and phonetic components occur. Phonetic components have a much more tightly constrained distribution than radicals in all cases.

Figure 6.

Analyses of the internal representations that develop as the model learns to read. Panel A depicts change in the mean Euclidean distance for comparisons among item types over training. Between P. Family = comparisons between all items that share the critical phonetic component (ㄅ) and control phonograms; Within P. Family = comparisons among all items that share the critical phonetic component; Orth. Control = comparisons between all items with the critical phonetic component and control items selected to share the same amount of orthographic information; Phon. Control = comparisons between all items with the critical phonetic component and their homophones that do not overlap orthographically. Panel B depicts the similarity space based on orthographic inputs for the test items. Grey patches indicate clusters of items that share orthographic structure; black circles indicate items that share a phonetic component. In Panel C, the similarity space based on hidden unit activations before training is shown. Grey patches and black circles as in Panel B. Panel D shows the similarity space based on hidden unit activations after 3 million trials of training on spelling to sound translation. The gray patch indicates items that share the same phonetic component, and cluster together only after training on spelling-to-sound correspondences. Black circles encompass items that share the same pronunciation in addition to sharing a phonetic component.

Figure 7.

Hidden unit activations for seven of the characters from the similarity analyses. Items from the same phonetic family as the sample item (苔/ $t^h a i$ /) have overlapping representations, whereas control items, matched for the degree of orthographic similarity, do not.

Table 1

Arrangement of radicals in complex characters

		Radical Sequence						
		1 st	2 nd	3 rd	4 th	5 th	6 th	7 th
乔	Right only					一	大	丩
侨	Left-right	亻				一	大	丩
矫	Left-right	丩	大			一	大	丩
莽	Top-bottom	艹				一	大	丩
臼	Center				臼			

Table 2

Materials for Study 1

	High frequency			Low frequency		
	R-C	R-I	I-I	R-C	R-I	I-I
Frequency	466.25	470.03	489.78	12.36	13.03	12.31
Consistency level	1.00	0.47	0.40	1.00	0.31	0.25
Family size	5.80	9.35	9.45	5.60	9.25	8.10
Number of radicals	3.10	2.80	3.05	3.50	3.30	3.25
Number of strokes	9.60	8.45	9.85	11.60	11.65	10.85

Note: Frequency expressed as occurrences per million

Table 3

Results of the ANOVA of SSE and accuracy of connectionist model in Study 1

	F1			F2			minF		
	df	f	p	df	f	p	df	f	p
SSE									
Frequency	1,39	928.39	<0.01	1,114	63.13	<0.01	1,128	59.11	<0.01
Type	2,78	253.61	<0.01	2,114	13.04	<0.01	2,126	12.40	<0.01
Freq × Type	2,78	158.44	<0.01	2,114	7.35	<0.01	2,124	7.024	<0.01
High Freq: Type	2,78	179.72	<0.01	2,115	0.28	0.757	2,115	0.27	0.757
Low Freq: Type	2,78	210.06	<0.01	2,115	12.96	<0.01	2,129	12.20	<0.01
Accuracy									
Frequency	1,39	238.00	<0.01	1,114	27.88	<0.01	1,137	24.95	<0.01
Type	2,78	58.27	<0.01	2,114	5.55	<0.01	2,135	5.067	<0.01
Freq × Type	2,78	58.27	<0.01	2,114	5.54	<0.01	2,135	5.059	<0.01
High Freq: Type	2,78	3.16	0.048	2,115	0.01	0.992	2,116	0.00	0.991
Low Freq: Type	2,78	59.26	<0.01	2,115	8.98	<0.01	2,148	7.798	<0.01

Table 4
Results of the ANOVA of RTs and accuracy of behavioral experiment in Study 1

	F1			F2			minF		
	df	f	p	df	f	p	df	f	p
RTs									
Frequency	1,35	174.15	<0.01	1,114	116.26	<0.01	1,129	69.71	<0.01
Type	2,70	52.45	<0.01	2,114	16.29	<0.01	2,169	12.42	<0.01
Freq × Type	2,70	59.82	<0.01	2,114	12.55	<0.01	2,156	10.37	<0.01
High Freq: Type	2,70	2.53	0.087	2,115	0.13	0.876	2,127	0.12	0.884
Low Freq: Type	2,70	74.74	<0.01	2,115	14.27	<0.01	2,154	11.90	<0.01
Accuracy									
Frequency	1,35	67.01	<0.01	1,114	30.39	<0.01	1,144	20.90	<0.01
Type	2,70	30.74	<0.01	2,114	15.96	<0.01	2,183	10.50	<0.01
Freq × Type	2,70	23.78	<0.01	2,114	8.39	<0.01	2,173	6.201	<0.01
High Freq: Type	2,70	2.18	0.121	2,115	0.50	0.609	2,160	0.40	0.667
Low Freq: Type	2,70	40.63	<0.01	2,115	18.90	<0.01	2,182	12.80	<0.01

Table 5

The interaction between frequency (times/million) and consistency in training corpus

	Frequency			All
	Low(≤ 10)	Middle(10–100)	High(≥ 100)	
Characters	1750	1537	1182	4468
Phonogram	1525(87%)	1239(81%)	651(55%)	3415(76%)
Number of Phonogram				
R-C	392 (26%)	258 (21%)	100 (15%)	750 (22%)
R-I	685 (45%)	553 (45%)	314 (48%)	1552 (46%)
I-I	422 (28%)	414 (33%)	229 (35%)	1065 (32%)
Consistency Level				
R-C	1.00	1.00	1.00	1.00
R-I	0.48	0.48	0.60	0.50
I-I	0.15	0.22	0.43	0.24

Table 6

Materials for Study 2

	Middle Frequency		Low Frequency	
	Consistent	Inconsistent	Consistent	Inconsistent
Frequency 1	45.06	38.72	3.48	3.13
Frequency 2	42.64	45.98	5.21	3.75
Consistent level	1.00	0.35	1.00	0.32
Number of Friends	6.03	3.17	5.97	3.4
Number of Enemies	0	3.03	0	2.93
Friends frequency	1646	593	1420	427
Enemies frequency	0	1125	0	1098
Number of radicals	3.27	3.03	3.33	3.2
Number of strokes	10.53	10.23	11.23	10.83
Phonetic frequency	692	627	728	657

Note: Frequency 1=frequency (per million) from *Modern Chinese Frequency Dictionary* (1986); Frequency 2=frequency (per million) from *Balanced Corpus of Modern Chinese* (Sun, 2006); the value of frequency was times per million.

Table 7

Results of the ANOVA of SSE and accuracy of connectionist model in Study 2

	F1			F2			minF		
	df	f	p	df	f	p	df	f	p
SSE									
Frequency	1,39	1313.56	<0.01	1,116	36.38	<0.01	1,122	35.3	<0.01
Consistency	1,39	120.37	<0.01	1,116	6.62	0.011	1,128	6.36	0.013
Freq $\times$ Con	1,39	89.48	<0.01	1,116	4.27	0.041	1,126	4.07	0.046
Mid-freq: Con	1,39	48.34	<0.01	1,117	0.10	0.754	1,117	0.09	0.753
Low-freq: Con	1,39	107.48	<0.01	1,117	8.27	<0.01	1,133	7.67	<0.01
Accuracy									
Frequency	1,39	703.44	<0.01	1,116	40.63	<0.01	1,129	38.4	<0.01
Consistency	1,39	66.78	<0.01	1,116	4.19	0.043	1,129	3.94	0.049
Freq $\times$ Con	1,39	41.53	<0.01	1,116	3.16	0.078	1,132	2.93	0.089
Mid-freq: Con	1,39	5.27	0.027	1,117	0.03	0.869	1,118	0.02	0.863
Low-freq: Con	1,39	55.97	<0.01	1,117	5.44	0.021	1,137	4.95	0.028

Table 8
Results of the ANOVA of RTs and accuracy of behavioral experiment in Study 2

	F1			F2			minF		
	df	f	p	df	f	p	df	f	p
RTs									
Frequency	1,39	139.61	<0.01	1,116	35.96	<0.01	1,153	28.5	<0.01
Consistency	1,39	78.55	<0.01	1,116	15.37	<0.01	1,149	12.8	<0.01
Freq × Con	1,39	33.94	<0.01	1,116	4.74	0.031	1,142	4.15	0.043
Mid-freq: Con	1,39	14.29	<0.01	1,117	1.17	0.282	1,134	1.08	0.300
Low-freq: Con	1,39	90.33	<0.01	1,117	14.32	<0.01	1,146	12.3	<0.01
Accuracy									
Frequency	1,39	76.91	<0.01	1,116	6.47	0.012	1,134	5.96	0.016
Consistency	1,39	237.89	<0.01	1,116	9.93	<0.01	1,125	9.53	0.003
Freq × Con	1,39	87.19	<0.01	1,116	6.47	0.012	1,132	6.02	0.016
Mid-freq: Con	1,39	5.21	0.028	1,117	0.18	0.675	1,125	0.17	0.678
Low-freq: Con	1,39	216.68	<0.01	1,117	15.49	<0.01	1,132	14.4	<0.01