

Variable Vulnerability of Words to Visual Impairment: Exploring Grain-Size Effects in Reading

Giovanni Pagliuca (gp517@york.ac.uk)
University of York, UK

Padraic Monaghan (p.monaghan@lancaster.ac.uk)
Lancaster University, UK

Rob McIntosh (r.d.mcintosh@ed.ac.uk)
University of Edinburgh, UK

Abstract

Languages vary in the grain size used to map orthography onto phonology, but the use of different grain sizes varies also within a language. A core question is what kind of information the reading system is using to map orthography onto phonology. A unit that has received particular attention in recent years is the grapheme. Graphemes are defined as written representations of phonemes, and they can be composed of multiple letters. If graphemes are functional reading units, then a word containing a multi-letter grapheme (digraph) should be read more accurately than a word without, given the same kind of perceptual noise to impair the orthographic input. We trained a neural network model to map English orthography onto phonology. Impairing the model revealed a distinction in performance with words with digraphs less affected by damage than control words. A behavioral experiment confirmed the prediction of the model. We discuss these results in light of a psycholinguistic grain size theory of reading.

Keywords: grain size; reading; modeling

Introduction

The question of how we recognize words is one of the core issues in cognitive psychology, and has been investigated at length (for a review see Prinzmetal, 1995). There are two possibilities for how written words may be identified: whole word recognition (or holistic recognition) approaches and recognition by parts approaches. In the former approach, a word is hypothesized to be recognized as a single object, as suggested by word-superiority effects in reading (Wheeler, 1970). The latter approach has instead tried to identify smaller grain sizes as important for word recognition such that some combinations of letters and their features (Rastle & Coltheart, 1998; Martelli, Majaj, & Pelli, 2005), orthographic syllables (Carreiras, Alvarez, & De Vega, 1993), and/or onsets and rimes (Bowey, 1990) are the functional units of the reading system. The debate has sparked great interest especially for its developmental and pedagogical implications (Ziegler & Goswami, 2005; Stanovich & Stanovich, 1995).

A main theoretical framework has been recently proposed which assumes that learning to read is learning to find shared grain sizes in orthography and in phonology (Ziegler

& Goswami, 2005), grain sizes that are language specific and allow for an efficient mapping between the two levels or representation. In this view the ability to identify and make optimal use of orthographic clusters in reading plays a major role in defining the ability to read.

One high-order cluster that has received particular attention in recent years as a potentially important grain size in reading is the grapheme. A grapheme is defined as a written representation of a phoneme (Berndt, Reggia, & Mitchum, 1987). For example, the word BREAD contains the five letters B, R, E, A, D, that map onto the four phonemes in /bred/. Therefore, four graphemes can be identified: B, R, EA, D, in virtue of the mapping between the letters and the phonemes. Graphemes therefore provide a direct mapping between orthography and phonology, and this mapping depends on the grain size of the orthographic representation.

Two main studies suggested explicitly that graphemes might act as perceptual units in reading. Rey, Jacobs, Schmidt-Weigand and Ziegler (1998), found that word identification times were longer for low frequency words composed of few graphemes (i.e. /biJ/ as in BEACH) than for control words composed of many graphemes (i.e. /krIsp/ as in CRISP), but critically having the same number of letters. This graphemic effect slowed processing time due to the extra-cost of having to functionally group more than one letter into a single grapheme. The study thus offered empirical evidence of a “phonemic” effect in orthographic processing.

A second study conducted by Rey, Ziegler and Jacobs (2000) implemented a letter detection paradigm to explore graphemic effects in reading. They asked participants to detect a single letter in words where the target letter was embedded in a multi-letter grapheme (i.e. ‘A’ in ‘BEACH’) or when it corresponded to a single-letter grapheme (i.e. ‘A’ in ‘PLACE’). Participants were slower in detecting the target in the former case than in the latter, suggesting that graphemes are processed as perceptual units in reading. This effect was found irrespective of the word frequency or phonemic similarity, and held for both English and French native speakers. The authors claim that the effect emerges from the difficulty in splitting the unit (grapheme) into its constituent letters in order to perform the task. However, neither of these studies directly tested participants on a reading task, and the claim that graphemes are “perceptual

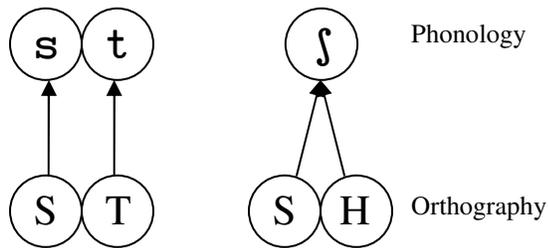


Figure 1. Different degrees of componentiality in the orthography to phonology mapping (left side: highly componential; right side: less componential).

reading” units remains inferential, with only indirect evidence supporting it. If graphemes are single perceptual reading units, then one prediction should follow: words containing multi-letter graphemes should be more resistant to impairment to the visual input and should be recognized more accurately than words without such graphemes. This should not be the case for single letters not part of multi-letters graphemes, that entail a highly componential mapping between orthography and phonology. Single-letter graphemes entail highly componential mappings (one-to-one), while multi-letter graphemes entail less componential ones (two-to-one) (see Figure 1).

The degree of componentiality in the mapping should therefore predict how well a particular grain size can be perceived.

Moreover, if the reading system develops sensitivity during the course of reading acquisition to letter clusters that functionally map onto one phoneme, then a neural network that employs single letters as input and single phonemes as output should develop the same sensitivity to words with graphemes, and show less vulnerability for words containing multi-letter graphemes in case of damage at the input level. In this view, graphemes would be considered as emergent functional units that the model develops naturally in virtue of the componential nature of the mapping between letters (orthography) and phonemes (phonology).

One way to explore this sensitivity to grain sizes is given by the possibility of impairing a standard model of reading and look at the level of accuracy with which different grain sizes are reproduced.

The scope of the present study is twofold. In the first part we look at whether a simple neural network with no explicit graphemic level of representation trained on a large corpus of English words could discover the existence of graphemes as an appropriate grain size for English.

We tested the model specifically on a set of words containing multi-letter graphemes (digraphs) and a set of control words, both with and without impairment to the model’s input. In the second part we confirm with a behavioral investigation the same prediction that words containing multi-letter graphemes are more resistant to visual noise than control words, testing healthy participant with degraded stimuli.

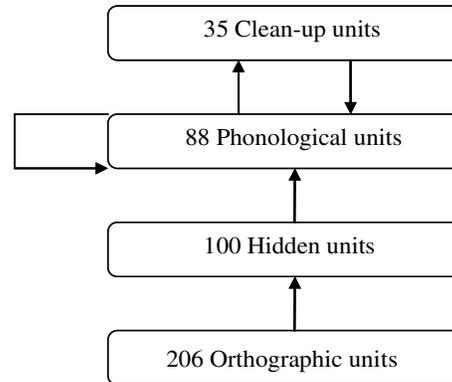


Figure 2. Architecture of the model.

Discovering grain-sizes in a model of reading

Computational models of reading have offered different views on the nature of the structure and the development of grain sizes. Dual Route Cascade models (DRC; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Perry, Ziegler, Zorzi, 2007) have explicitly implemented a level of representation for the graphemes as model’s input, instead of using single letters. Classic parallel distributed processing (PDP) models of reading have avoided implementing such levels of representation (Harm & Seidenberg, 1999; Harm & Seidenberg, 2004; Plaut, McClelland, Seidenberg & Patterson, 1996). This latter class of models generally use single letters as input and assume that information about the appropriate grain size would be discovered during the course of learning to map orthography onto phonology. This ability has been shown in the simulation of graded consistency effects that the DRC fails to account for (Zevin & Seidenberg, 2006). A PDP model then should show sensitivity to variations in grain sizes within a language. Following we describe an attempt to explore this sensitivity by impairing a standard model of reading.

Method

Architecture

The model is shown in Figure 2, and is based closely on Harm and Seidenberg’s (1999) connectionist model of reading. The network comprised of an orthographic layer with 206 units, a hidden layer with 100 units, and a phonological layer with 88 units. The orthographic units were fully connected to the hidden units, which in turn were fully connected to the phonological output units. A set of 35 cleanup units mediated the computation from the phoneme representation to itself (as in Harm & Seidenberg, 1999). The phonological units were self-connected with each other. The direct connections between phonological units allowed the encoding of simple types of dependencies between phonological features (the fact for example that a phoneme cannot be both *consonantal* and *sonorant*; Harm & Seidenberg, 1999). The orthographic representations for words were slot based, as in Harm and Seidenberg (1999)

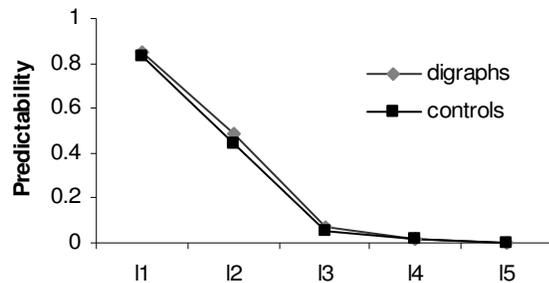


Figure 3. Matching in levels of predictability per letter position for words with digraphs and control words.

with one of 26 units active in each slot for each letter. There were 8 letter slots. Words were inputted to the model with the first vowel of the word in the fourth letter slot. At the output layer, there were 8 phoneme slots for each word, three representing the onset, two for the nucleus, and three for the coda of each word. So, the word “help” was inputted as “- - h e l p - -”, where “-” indicates an empty slot. This kind of representation has the advantage of capturing the fact that different phonemes in different positions sometimes differ phonetically (Harm & Seidenberg, 1999), though similarities between onset and coda phonemes were not directly available to the model. Each phoneme was represented in terms of 11 phonological features, as employed by Harm and Seidenberg (1999). The units in the model were standard sigmoidal units with real-valued output ranging between 0 and 1 for the input and hidden layer, and -1 and 1 for the output layer.

Training and testing

7291 monosyllabic words were chosen from the CELEX database (Baayen, Pipenbrock, & Gulikers, 1995) and were used as the training corpus. We selected words with frequency greater than 68 per million in the database. Each word was 1 to 8 letters long and was assigned a frequency derived from its frequency of representation in the original database and log transformed (see Plaut et al., 1996, for a discussion of log frequency compression). Words with more than three phonemes in the coda were omitted from the input set. The model was trained using the backpropagation learning algorithm (Rumelhart, 1986), with the weight of the connections initialized to random values with mean 0, variance 0.5, and learning rate $\mu=0.001$.

Words were selected randomly according to their frequency. We stopped training after ten million words had been presented. For testing, we assessed the model’s performance for a reading task on all words in the training set.

The model’s production for each phoneme slot at the output was compared to all the possible phonemes in the training set, and to the empty phoneme slot (where values of all units in the slot were -1). For word presentations, if the model’s performance in each phoneme slot was closer to the target than to any other phoneme, then the model was judged to have read the word correctly.

The model was specifically tested on two sets of 62 words each. One set contained words with digraphs in the initial position and one did not. All the stimuli were monosyllabic, 5 letters long and occurred in the training set.

Three types of digraphs were identified, with the constraint of being in initial position: SH (as in *shame*), CH (as in *chair*), and TH (as in *thorn*). All these three digraphs map respectively onto one phoneme slot in the output (SH $\rightarrow \int$; TH $\rightarrow \theta$ or δ ; CH $\rightarrow t\int$) as affricates were represented as a single set of phonological features in the output.

The control words were matched with the experimental ones for: same initial letter, same number of letters, similar bigram frequency for the first two bigrams, and word frequency. The control words for SH digraph words all began with ST, for CH, the control words began with CR, and for the TH words, control words began TR. The stimuli were also matched for likelihood of guessing the word using a procedure to estimate predictability of words from partial views (see Figure 3). For each letter in the word we estimated the likelihood of correctly reading the word if that letter was the leftmost letter processed (and all rightward letters were processed correctly). The estimate was derived by dividing the lexical frequency of the target word by the summed lexical frequency of all words between 3 and 9 letters in length that would be consistent with that partial template of the word.

Input Impairment

In order to impair the visual input to the model across the critical letters, we reduced the activation from input letter slots along a monotonic gradient from left to right, such that the largest reduction in activation was from the leftmost letter slots. We used two severities of impairment. A severe impairment was implemented by reducing the weights of the connections to the leftmost (first) letter slot by 75%, the connections to the second letter slot by 64%, reduction of 53% to the third letter slot, and so on across the input units. Hence, the gradient was linear, and ended with a 0% reduction of connections to the rightmost unit. A mild impairment reduced connections from the first letter slot by 50%, the second slot was reduced by 43%, the third slot by 36%, and so on, to zero reduction for the rightmost letter slot. Both words with digraphs and their control occupied the same number of slots at the input, starting from the same leftmost slot which contained the first letter. This type of impairment procedure ensured that the beginning of the word (the first leftmost letter slots, which were occupied by digraphs and control digrams in our experimental material) was most affected by the impairment to the input.

Results

After it had been trained on 10 million patterns, the model correctly reproduced 93.6% of the words in the corpus, which falls slightly short of the performance achieved by the Harm and Seidenberg (1999) model of reading (98% correct). Before impairment, the model was 100% correct at reading both experimental stimuli and control words. The model with the severe input impairment made errors on reading 83.5% of the words. Of these, 75% of the word-reading errors affected only the first and/or second phonemes. Two types of errors were measured: omission of the first and/or second phoneme and substitutions of the first

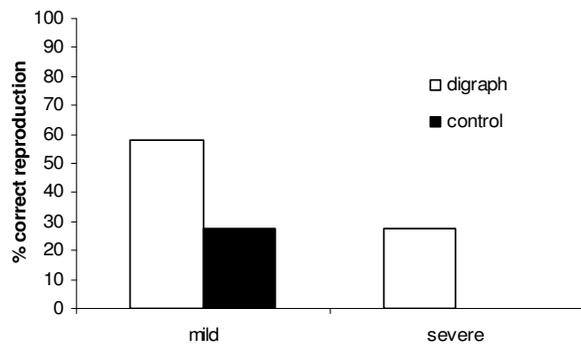


Figure 4. Levels of accuracy for words with digraphs and control words for mild impairment and severe impairment.

and/or second phoneme, or a combination of the two. The mild impairment model made errors on 56.3% of the words. Errors on the first 2 phonemes accounted for 66.2% of the errors. The impairment procedure implemented at the input level was thus successful in selectively impairing the output at its leftmost side.

The model was then assessed on levels of accuracy for the set of experimental words and their controls. Both the severe and mild impairment model showed sensitivity to the presence of digraphs in a word: words with digraphs were reproduced more correctly than words without ($\chi^2(1) = 11.896, p < .01$). The same sensitivity was again shown by the mild impairment model ($\chi^2(1) = 19.701, p < .01$) (see Figure 4). Together, both models showed a relative robustness to impairment for words containing digraphs when compared to control words with no digraphs.

Discussion

The simulation presented above confirmed our prediction that a neural network trained to map letters onto phonemes performs better with words containing digraphs (or multi-letter graphemes) than with control words which do not contain digraphs, when there is impairment to the model's input. The model does so as a consequence of the fact that the two letters embedded in the digraphs map onto a single phoneme, and although both letters are damaged, the combined residual activation of the two letters is still enough to activate the right phoneme at the output level and preserve the whole word. The model does not have any explicit level of representation for the graphemes, but employs only single letter slots at the input (orthographic) level. The graphemic effect shown here is therefore a consequence of the initial phoneme in the digraph words receiving partial information from more than one letter position, whereas the control words required a one-to-one mapping from letter slots to phoneme slots. In the next study, we tested whether the model's demonstration of different grain-size effects for words containing digraphs was evident in a behavioral study, when visual representation of the word was impaired for participants reading words.



Figure 5. Example of stimulus used in the behavioral study. The word "chain" is here represented.

Behavioral Experiment

In this experiment we directly tested the grain-size effect in a reading task. In contrast to the model's performance, participants tend to guess the identity of a word from partial visual information.

Method

Participants. Fifteen university students, all native English speakers, participated in the experiment. All had normal or corrected to normal vision.

Materials. The same sets of words used in the simulation were used in the behavioral study. In addition to the experimental stimuli and control words 84 fillers were selected, all five letters long. These fillers had different initial bigrams and different initial letters to the experimental and control stimuli.

Words were presented centered in a frame with the addition of random 2-dimensional digital pixel noise applied according to a gradient from left to right, left being the most affected by digital noise (example is shown in Figure 5), in order to instantiate a similar degraded impairment to the visual information as employed in the simulation work. This procedure made the initial letters difficult to perceive, leaving the rightmost letters almost unaffected. The same gradient and amplitude of digital noise was applied to all the stimuli.

The experimental session consisted of 3 blocks. The first two blocks contained 20 experimental stimuli with digraphs, 20 control stimuli without digraphs, and 28 fillers, and the third block contained the remaining 22 experimental stimuli, 22 control words, and 28 fillers. The order of the blocks was randomised across participants, as were the stimuli in each block.

Procedure. Participants were asked to read aloud the stimuli. They were informed that all the stimuli were real words and that they should report whatever they could read. Participants saw a fixation cross at the center of a computer screen for 1000-ms, followed by the word in a frame with the added pixel noise for 250-ms. The stimuli were then followed by a mask for 500-ms. An interstimulus interval of 2000-ms preceded the presentation of the next fixation cross. Responses were recorded and subsequently transcribed by the experimenter. Participants were seated approximately 25cm from the screen, and stimuli were presented centered as heavy gray text (Courier New 150) on a gray scale background (300 by 100 px) in order to

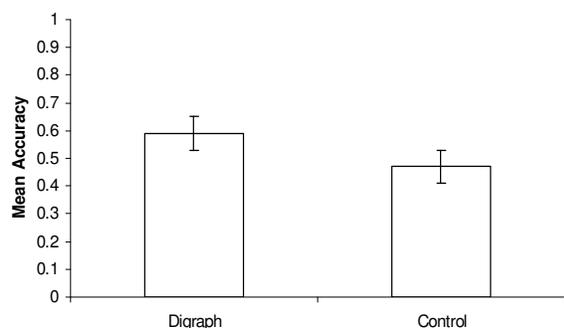


Figure 6. Means of accuracy for words with digraphs and words without.

reduce contrast. A brief practice session preceded the experiment.

Results

Participants made errors that affected exclusively the first and/or second phonemes of the words for a mean of 42% of the responses. The same scoring procedure used for the model was implemented to score the participants responses. The digital noise added to the frame was therefore successful in impairing specifically the beginning (left side) of the words.

A Linear Mixed Effects Analysis (Baayen, 2007) was performed with the factor “digraph” as a fixed effect and words and participants as random effects. A word was judged to be read accurately only if all its phonemes were pronounced correctly. A main effect of digraph status was found, with words with digraphs being reported more accurately than words without, $F(1, 107.998) = 17.177$, $p < .01$ (see Figure 6).

Discussion

The results from the behavioral experiment confirmed the prediction that words containing digraphs are read more accurately than control words without when participants view words with noise that impairs accurate perception of the stimuli. The effect cannot be accounted for by the use of guessing strategies, nor by variables such as length or frequency (at the word level or bigram level), as those variables were controlled across the two groups of stimuli used. The experiment provides direct evidence that the size of the graphemes in the word have a differential effect on reading, indicating different grain size effects within a language during reading.

General Discussion

The present study explored grain size effects in reading using a computational model of reading single words in English, and a behavioral study testing a prediction raised by the model’s performance. Ziegler and Goswami’s (2005) theoretical model of different grain size effects in different orthographies proposed that the extent to which letters mapped regularly and compositionally onto phonemes in

words determined the type of processing that occurs in the reading system. For shallow orthographies, such as Italian or Spanish, where the majority of letter-phoneme mappings are one-to-one, the grain size developed during reading is small. For languages where the context of individual letters has to be taken into account to pronounce them, then the grain size grows according to the size of the context required. English has instances both of componential, one-to-one mappings between letters and phonemes, and also occasions where sets of letters (digraphs, or trigraphs) map onto single phonemes. Consequently, we investigated whether, within a language, the grain size for processing was sensitive to the local characteristics of the word in terms of componentiality.

The model of reading suggested that the computational properties of the mapping between letters and phonemes was differentially vulnerable to variations in the mapping within individual words. Though the difference between words with digraphs and words without was not observable in the unimpaired model, when we affected the veridicality of the input to the model, differences between these word types was seen. Digraphs were less vulnerable to impairment than were single letters that mapped to phonemes.

This was because the model was able to take lowered activation from both letters of the digraph to produce a pronunciation for that phoneme. In contrast, the words without digraphs had to be processed in a componential way, such that an impairment to the activation of a single letter had a profound effect on the model’s ability to process the pronunciation dependent on that letter.

Alternative traditions of modeling reading vary in the extent to which they explicitly encode graphemes in words. Our model did not specify the graphemes in the language, yet we showed that a parallel distributed processing model of reading can discover different grain sizes in the course of learning to map single letters onto phonemes, and that therefore an explicit level of representation for the graphemes at the input level might not be needed. The model developed sensitivity to the graphemes as a consequence of the regularities in the lexicon in mapping letters and sets of letters onto phonemes.

We then applied the prediction raised by the model to investigate whether graphemes might have a special status in reading in adults, as suggested by previous studies (Rey, Jacobs, Schmidt-Weigand & Ziegler, 1998; Rey, Ziegler and Jacobs, 2000). These studies did not test participants on a reading task directly, but instead used either word identification or a letter detection task to test the special status of graphemes. Although consistent with the main assumption that graphemes are perceptual units in reading, these studies did not directly demonstrate that digraphs had a differential effect on word naming tasks.

In the behavioral study we confirmed the prediction made by the model that words with digraphs are more robust to an impairment to the visual input for reading. When we controlled sets of words for length, neighbours, frequency, and predictability of the word, we found that, when pixelated noise was added to the left of visually presented words, the words containing digraphs were read more

accurately than words with one-to-one letter-phoneme mappings.

Taken together, the modeling work determined that different computational properties of the mapping between letters and phonemes suggested that psycholinguistic effects of words should vary according to the compositionality of the mapping. The behavioral study confirmed that these computational properties do have a profound effect on reading, at least under conditions where the visual input is impaired. The results are consistent with the hypothesis that graphemes have a special status in reading, due to their computational properties, and that they emerge in the course of a system learning the regularities between orthographic and phonological representations of words.

Acknowledgments

This research was supported by an EU Sixth Framework Marie Curie Research Training Network Program in Language and Brain: <http://www.hull.ac.uk/RTN-LAB/>. We are grateful to Martin Corley and Tony Morland for the development of some of the scripts used in this paper, Andy Ellis and Tom Hartley for helpful critical comments.

References

- Baayen, R.H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX Lexical Database* (Release 2) [CD-ROM]. Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.
- Baayen, R.H. (2007). *Analyzing linguistic data: A practical introduction to statistics*. Cambridge: Cambridge University Press.
- Berndt, R.S., Reggia, J.A., & Mitchum, C.C. (1987). Empirically derived probabilities for grapheme-to-phoneme correspondences in English. *Behavior Research Methods, Instruments & Computers*, 19, 1-9.
- Bowey, J.A. (1990). Orthographic onsets and rimes as functional units of reading. *Memory & Cognition*, 18, 419-427.
- Carreiras, M., Alvarez, C.J., & De Vega, M. (1993). Syllable frequency and visual words recognition in Spanish. *Journal of Memory and Language*, 32, 766-780.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J.C. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204-256.
- Harm, M.W., & Seidenberg, M.S. (1999). Phonology, reading and dyslexia: insights from connectionist models. *Psychological Review*, 163, 491-528.
- Harm, M.W., & Seidenberg, M.S. (2004). Computing the meaning of words in reading: cooperative division of labor between visual and phonological processes. *Psychological Review*, 111, 662-720.
- Martelli, M, Majaj, N.J, & Pelli, D.G. (2005). Are faces processed like words? A diagnostic test for recognition by parts. *Journal of Vision*, 5, 58-70.
- Perry, C., Ziegler, J.C., Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review*, 114, 273-315.
- Plaut, D.C., McClelland, J.L., Seidenberg, M.S., & Patterson, K.E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56-115.
- Prinzmetal, W. (1995). Visual feature integration in a world of objects. *Current Directions in Psychological Science*, 4, 90-94.
- Rastle, K., & Coltheart, M. (1998). Whammy and double whammy: length effect in nonword naming. *Psychonomic Bulletin and Reviews*, 5, 277-282.
- Rey, A., Jacobs, A.M., Schmidt-Weigand, F., & Ziegler, J.C. (1998). A phoneme effect in visual word recognition. *Cognition*, 68, 41-50.
- Rey, A., Ziegler, J.C., & Jacobs, A.M. (2000). Graphemes re perceptual reading units. *Cognition*, 75, B1-B12.
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature*, 323, 533-536.
- Stanovich, K.E., Stanovich, P.J. (1995). How research might inform the debate about early reading acquisition. *Journal of Research in Reading*, 18, 87-105.
- Wheeler, D.D. (1970). Processes in word recognition. *Cognitive Psychology*, 1, 59-85.
- Zevin, J. D., & Seidenberg, M. S. (2006). Simulating consistency effects and individual differences in nonword naming. *Journal of Memory and Language*, 54, 145-160.
- Ziegler, J.C., & Goswami, U. (2005). Reading acquisition, developmental dyslexia, and skilled reading across languages: A psycholinguistic grain size theory. *Psychological Bulletin*, 131, 3-29.