

New Analyses of Lexical Influences on the Processing of Pseudo-homophones in the Lexical Decision Task: Still More Challenges for Models of Visual Word Recognition

Robert S. McCann (rmccann@uwaterloo.ca)

Department of Psychology, 200 University Avenue
Waterloo, Ontario, Canada N2L 1P8

Blair C. Armstrong (blair.armstrong@utoronto.ca)

Department of Psychology, 1265 Military Trail
Toronto, Ontario, Canada M1C 1A4

Michael Reynolds (michaelchanreynolds@trentu.ca)

Department of Psychology, 1600 Westbank Drive
Peterborough, Ontario, Canada K9J 7B8

Derek Besner (dbesner@uwaterloo.ca)

Department of Psychology, 200 University Avenue
Waterloo, Ontario, Canada N2L 1P8

Abstract

New analyses of pseudo-homophone RTs (e.g., BRANE) from two published lexical decision studies clarify lexical involvement in pseudo-homophone processing and challenge widespread assumptions about word frequency effects. First, RTs increased along with increases in the proportion of base-word letters that appeared in the pseudo-homophone (e.g., WHELT-WELT slower than PHAWT – FOUGHT) suggesting that “No” decision-making is slowed by mutually reinforcing activation in phonological and orthographic representations of base word knowledge. Second, effects of base-word frequency were either extremely weak or nonexistent among pseudo-homophones that contained most or all the letters that make up their base word. In contrast, among pseudo-homophones that shared fewer letters with their base word (e.g., “PHAWT”), RTs for items derived from high-frequency base words were *faster* than RTs for items derived from low-frequency base words. These findings (i) challenge the ubiquitous assumption that lexical representations are frequency sensitive and (ii) suggest that lexical decision involves a spell-check.

Keywords: Pseudo-homophones; Orthographic Similarity; Lexical Decision; Word-Frequency Effects

Introduction

Investigators of visual word recognition have long sought to explain how a computational system that’s almost exclusively the product of processing words is able to rapidly and accurately discriminate words from orthographically legal and pronounceable nonwords. Important clues to this ability have emerged from investigations of performance on pseudo-homophones (PH’s), which are nonwords at the level of spelling (e.g., “BRANE”), but whose pronunciation matches that of a known word (e.g., BRAIN). In the lexical decision task, “No” responses to PHs are typically slower and more error prone than responses to matched nonhomophonic

nonwords (Besner & Davelaar, 1983; Coltheart, Davelaar, Jonasson, & Besner, 1977; McCann, Besner, & Davelaar, 1988; Ziegler, Jacobs, & Kluppel, 2001). These “PH effects” have been simulated in both localist and PDP models of visual word processing (e.g., Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Harm & Seidenberg, 2004).

Early accounts of PH effects in lexical decision held that a PH activates the representation for the PH’s base word (BW) in a phonological lexicon (Coltheart et al., 1977; Besner & Davelaar, 1983). Activation of the BW’s phonological representation interferes with the correct “No” response because activation of lexical representations provides evidence that the item under consideration IS a word.

A ubiquitous assumption in the word recognition literature is that lexical representations are frequency sensitive (Morton, 1969). For example, the broadly influential interactive-activation model (McClelland & Rumelhart, 1981) assumes that lexical representations for high-frequency words possess higher levels of resting activation than representations for low-frequency words. Functionally, the difference in resting activation makes lexical access faster for high frequency than for low-frequency words which, in turn, yields performance advantages for high-frequency words in tasks such as lexical decision.

If PH processing is influenced by contact with frequency-sensitive phonological representations of the BW, PH performance should show effects of BW frequency (BWF). However, in their lexical decision study, McCann, Besner, and Davelaar (1988; hereafter, MBD88) found this not to be the case; the correlation between PH RTs and BWF was a mere .058, $p > 0.5$.

MBD88 argued that the null result of BW frequency challenged the received view that lexical representations are inherently frequency sensitive.

MBD88's findings notwithstanding, the widespread assumption that frequency-sensitive lexical representations affect PH processing is as entrenched today as it was in the 1980's. For example, Coltheart and colleagues' Dual-Route Cascaded (DRC) model (Coltheart et al., 2001) assumes that PH processing involves competitive/cooperative interactions between two frequency-sensitive lexical representations: The BW entry in the orthographic input lexicon, which represents knowledge about BW spelling, and the BW entry in the phonological output lexicon, which represents knowledge about BW phonology. These interactions are more vigorous, and result in higher levels of mutual activation, for high-frequency BWs than for low-frequency BWs (Ziegler, Jacobs, & Kluppel, 2001).

In parallel distributed processing (PDP) models such as those based on the triangle framework (Harm & Seidenberg, 2004), word frequency is encoded in the strengths of the connections between elementary processing units arranged in networks that represent word-specific orthographic, phonological, and semantic knowledge in the form of patterns of activation across these units. The connection strengths, established through training regimens that sample words based on their printed frequency, cause the units to settle into word-specific patterns - "attractor points" - more rapidly for high- than for low-frequency words. That, in turn, influences how strongly and effectively activation propagates from a set of units representing one form of knowledge to the others. PH processing in the lexical decision task is slowed by broader levels of activation across the units of the semantic network, relative to nonwords that are not homophonic with any known word (Harm & Seidenberg, 2004).

Both distributed and localist computational models would appear to predict that RTs to PHs that sound identical to high-frequency BWs should be slower than RTs to PHs that sound identical to low-frequency BWs. Intriguingly, both classes of models also hold that PH processing is sensitive to orthographic relations between the PH and its BW. We have already noted that in localist models of word processing such as DRC, activation of lexical entries is strongly influenced by interactions between corresponding entries in the orthographic and phonological lexicons. The stronger a PH activates the orthographic lexical representation for its BW, the more the orthographic and phonological lexical representations for the BW will resonate, and the stronger the overall level of lexical activity will become. However, the extent to which a PH activates its BW representation in the orthographic input lexicon is influenced by the proportion of the letters that the BW shares with its PH. If the BW and PH share few (or no) letters (e.g., PHAUT-FOUGHT) the BW entry in the orthographic lexicon is activated only

weakly, damping the orthographic/phonological feedback loop, and minimizing the lexical activation "boost" that accrues to the PH (Coltheart et al., 2001). Therefore, in lexical decision, RTs to PHs that share few or no letters with their BW are *faster* than RTs to PHs that share all or most of their spelling with the base word.

For PDP models, PH processing is influenced by the extent to which the PH generates activity in the units that represent word-specific *semantic* knowledge. Harm and Seidenberg's (2004) simulations using a PDP model showed that PH's that share most of their letters with their base word activate a broader set of semantic features in the semantic network than PH's that share fewer letters. Consequently, in lexical decision, RTs to high letter-overlap PHs should be *slower* than RTs to lower-overlap PH's.

In summary, both localist and PDP models assert that "No" responses in lexical decision should *increase* (1) as the orthographic similarity between the PH and BW increases and (2) as the frequency of the BW increases. In addition, given that the two variables influence common processes, an interaction would also be expected, such that the more spelling overlap there is between the BW and the PH, the stronger the effects of BWF should be. Specifically, RTs should be slowest among PH's that share most of their letters with a high-frequency BW.

How do these predictions fare against published findings? Contrary to predictions, RTs to PHs have not been found to increase with BWF. MBD88 reported a null effect of BWF on PH RTs. Moreover, subsequent studies have reported that RTs for PHs derived from high-frequency BWs are *faster* than RTs for PHs derived from low-frequency BWs (Van Orden, 1991; Van Orden, Stone, Garlington, Markson, Pinnt, Simonfy, & Bricchetto, 1992; Ziegler, Jacobs, & Kluppel, 2001). As Ziegler, Jacobs, and Kluppel (2001) and Harms and Seidenberg (2004) both noted, the effects of BWF are *opposite* to model predictions.

As for the models' prediction of an impact of BW letter overlap with its PH, preliminary results are not encouraging. We correlated the MBD88 PH RTs with two standard orthographic similarity metrics, Levenshtein's (1966) Orthographic Distance (OD) measure and Van Orden's (1987) modification to Weber's (1970) graphical similarity measure. Neither OD ($r = .15, p = .17$) nor Van Orden's (1987) graphical similarity measure ($r = .14, p = .21$) accounted for significant variance.

Finally, to the best of our knowledge, no published study has tested for an interaction between the level of orthographic overlap between a PH and its BW and BWF.

Present Aims

The analyses reported here revisit the issues of whether (i) orthographic overlap between the PH and its BW, and (ii) BW frequency, influence lexical decision performance to PHs. Furthermore, (iii) we explicitly test for an interaction between letter-level overlap and BWF.

Collectively, these analyses provide several strong tests of model predictions.

Quantifying Orthographic Similarity. As we just noted, traditional measures of orthographic overlap between PHs and their BWs failed to account for significant variance in the MBD88 PH RTs. How might we reconcile these null results with model predictions? Perhaps, following Coltheart et al., (2001), the strength with which a PH activates lexical orthographic representations is determined primarily by the proportion of its letters that the BW shares with the PH. Both OD and Van Orden's graphical similarity measure are sensitive to orthographic dimensions that are orthogonal to simple shared letter proportion. For example, the values of both measures decrease when the PH is longer (contains more letters) than the BW (e.g., MIRTH-MIRTHE). Adding letters to create a PH does not alter the proportion of the letters in the BW that also appear in the PH, however. Thus, the traditional similarity metrics may incorporate irrelevant variance that reduces their power to predict PH RTs.

Accordingly, we derived a new orthographic similarity measure that quantifies *only* the proportion of the letters in the BW that also appear in the PH. This measure, SLC/lenBW, sums the number of letters that the BW shares with the PH (Shared Letter Count, abbreviated to SLC), and divides the SLC by the total number of letters in the BW (abbreviated to lenBW).

Following Balota, Yap, Cortese, Hutchison, Kessler, Loftis, Neely, Nelson, Simpson, and Treiman (2007), we quantified BWF using HAL-F, a log-transformed version of the HyperSpace Analog to Language (HAL) frequency counts (Lund & Burgess, 1996) obtained from Balota and colleagues' English Lexicon Project website (Balota et al., 2007). We then regressed SLC/lenBW, BWF, and their interaction on lexical decision PH RTs from MBD88 and Armstrong and Plaut (2016). To anticipate, the results from these two data sets were very similar and challenge predictions of both localist and PDP models.

Analyses of the MBD88 data set

A regression analysis was performed on the MBD88 PH RTs with mean PH RTs as the outcome variable and SLC/lenBW, BWF, and their interaction (centered) as the predictor variables. As SLC/lenBW increased, PH RTs also increased ($b = 219.4$, $t = 4.18$, $p < .001$). As shown in Figure 1, the zero-order relationship between SLC/lenBW and PH RTs was quite substantial. There was no main effect of BWF ($b = -3.25$, $t = 1.24$, $p > .2$), but critically, BWF participated in a significant SLC/lenBW by BWF interaction ($b = 85.79$, $t = 2.338$, $p < .03$).




Figure 1: Effect of BW-PH letter overlap (SLC/lenBW) on PH RTs in MBD88.

To illustrate the form of the interaction, we separated the 80 MBD88 PH's into those with SLC/lenBW values of 0.75 and above (high-orthographic-overlap items; $N = 66$) and those with SLC/lenBW values below 0.75 (lower-orthographic-overlap items; $N = 14$). The scatterplots in Figure 2 below illustrate the effects of BWF on RTs for the two categories of items. As shown in the left panel, BWF effects were strikingly absent among the PHs with the highest SLC/lenBW values ($r = 0.01$, $b = -0.25$, $t = -0.09$, $p = .93$). In sharp contrast, among the PHs with lower SLC/lenBW values, RTs to PHs constructed from high-frequency BWs were *faster* than RTs to PHs constructed from low-frequency BWs, though the main effect of BWF was only marginally significant ($r = 0.49$, $b = -7.94$, $t = -1.99$, $p < .08$).

Discussion. The performance patterns revealed by our new analyses appear inconsistent with computational models of word recognition and the lexical decision task. According to the word recognition models, activation of BW representation(s) should be strongest among PHs that contain most or all of the letters that make up their BWs, particularly for PHs derived from HF BWs. Assuming that "No" decision making is influenced by the strength with which the nonword activates lexical representations (Wagenmakers, Steyvers, Raaijmakers, Van Rijn, & Zeelenberg, 2004), PH RTs should be slowest among PHs that share the most letters with their BWs, particularly for PHs derived from HF BWs. Instead, RTs to PHs with the most letter-level overlap with their BWs, while indeed slower than RTs to PHs with less overlap, were *insensitive* to BWF. In contrast, among PHs with more distinct orthographies, PHs derived from high-frequency BWs were *faster* than RTs to PHs derived from low-frequency words.




Figure 2. Scatterplots of MBD88 Lexical Decision RTs as a function of BWF for PHs of higher (SLC/lenBW values of 0.75 and above; left panel) and lower (SLC/lenBW values < 0.75; right panel) levels of orthographic overlap with their BWs.

When surprising data patterns are revealed through post-hoc analyses, some skepticism is warranted regarding the robustness and generality of the findings. Furthermore, since only 14 of the 80 PHs in the MBD88 item set had SLC/lenBW values less than 0.75, critical aspects of the results are based on a very small sample of items. The small sample size raises the uncomfortable possibility that the critical items possess idiosyncratic attributes that might covary with the variables of interest, and so effects obtained with these items may not generalize to other PH samples.

MBD88's experiment was designed to study whether PH RTs would be slower than a set of nonhomophone controls, and sensitive to BWF. In contrast, a more recent lexical decision study by Armstrong and Plaut (2016) was designed to examine semantic ambiguity effects on word performance as a function of the makeup of the nonwords. Following Rodd, Gaskell, and Marslen-Wilson (2002)'s findings that significant semantic ambiguity effects emerge when the foils are PHs but not word-like nonwords, the nonwords in one of Armstrong and Plaut's conditions consisted of 500 PHs. No yoked nonword controls for the PHs were included in the study, and no analyses of PH RTs or error rates were reported. However, Armstrong and Plaut's PH set included many more PHs, and many more with SLC/lenBW values less than 0.75, than the MBD88 sample. Additionally, Armstrong and Plaut included a between-participants manipulation of stimulus quality by presenting the stimuli either in white against a dark background (high contrast condition) or in dark grey against a dark background (low contrast condition). Thus, the Armstrong and Plaut database of PH RTs provides multiple opportunities to test the robustness and generality of our findings with the MBD88 items, while also avoiding potential experimenter bias during item selection (Forster, 2000).

Analyses of Armstrong and Plaut (2016) data set

Prior to analyzing the Armstrong and Plaut (2016) PH RTs, we removed a large number of double PHs, that is, PHs that sound like two BWs with different spellings, as calculating unambiguous values for SLC/lenBW for double PHs is not possible. We also removed smaller numbers of PHs that we either couldn't associate immediately with a BW, or that Armstrong and Plaut classified as outliers. This left a sample of 363 PHs, 75 of which had SLC/lenBW values less than 0.75.

We then submitted the Armstrong and Plaut (2016) PH RTs to a multiple regression analysis including Stimulus Quality, SLC/lenBW, BWF, and the (centered) SLC/lenBW by BWF interaction as predictors. There was a significant effect of stimulus quality, with high contrast items being faster than low-contrast items ($b = 61.9$, $t = 17.9$, $p < .001$). More important for present purposes, the main effects of both SLC/lenBW ($b = 61.6$, $t = 4.68$, $p < .001$) and BWF ($b = -1.63$, $t = 2.12$, $p < .04$) were significant. Even more critically, we once again observed an interaction between SLC/lenBW and BWF ($b = 18.4$, $t = 3.54$, $p < .001$). The nature of the interaction is illustrated in Figure 3, which plots the effect of BWF for PHs with SLC/lenBW proportions of 0.75 or greater ($N = 288$) on the left side of the Figure, and PHs with SLC/lenBW proportions lower than 0.75 ($N = 75$) on the right side. The scatterplots for the high-contrast stimuli (top panels) reveal exactly the same pattern we found with the MBD88 PH set: No effect of BWF on RTs for PH's that share most or all of their letters with their base word ($t = 0.57$, $p > .5$), but a strong effect of BWF among PHs that share fewer letters ($t = -4.02$, $p < .001$). Again, the BWF effect reflected the fact that RTs to PHs derived from higher-frequency BWs were faster than RTs to PHs derived from lower-frequency BWs.

The two plots along the bottom of Figure 3 reveal exactly the same patterns when the PHs were perceptually degraded: A null effect of BWF on the high SLC/lenBW




Figure 3. Scatterplots of PH RTs as a function of BWF from Armstrong and Plaut (2016.) The plots on the left side of the Figure include RTs for PHs with high levels of letter overlap with their BWs. The plots on the right side of the Figure include RTs for lower-letter-overlap PHs. The upper panels plot the data from the high-contrast (clear) condition; the lower panels plot the data from the low-contrast (degraded) condition. Solid lines are the least-squares regression lines predicting the effect of BWF on each set of data. AP16 = Armstrong and Plaut (2016).

items ($t = 0.72$, $p > 0.4$), but faster RTs for PHs derived from high-frequency BWs than from low-frequency BWs among the lower SLC/lenBW items ($t = -3.19$, $p < .01$).

A reviewer of an earlier version of this paper was concerned that SLC/lenBW might be confounded with the extent to which a PH's spelling overlaps with English words in general, which could compromise the interpretation of the SLC/lenBW by BWF interaction. To address this concern, we performed a second multiple regression analysis including an additional predictor variable, OLD20, that quantifies the orthographic "typicality" of each PH's spelling pattern (Yarkoni, Balota, & Yap, 2008). The SLC/lenBW by BWF interaction was statistically unaffected by the inclusion of OLD20.

Discussion. The form of the SLC-SP/lenBW by BWF interaction in Armstrong and Plaut's (2016) PH data set replicates our findings with the MBD88 PH set in all respects and helps clear up inconsistent findings in the literature. Most notably, the lack of a statistically significant BWF effect in MBD88 conflicts with the results of several subsequent studies (Van Orden, 1991; Van

Orden et al., 1992; Ziegler, Jacobs, & Kluppel, 2001) where faster latencies were found for PH's derived from high-frequency BWs than from low-frequency BWs. The present analyses revealed faster latencies for PHs derived from high-frequency BWs in the MBD88 and Armstrong and Plaut (2016) data sets too, *but only* among PHs that are orthographically distinct from their BWs. Therefore, the absence of a main effect of BWF in MBD88 was almost certainly a consequence of the small number of low-spelling-overlap items in the MBD88 PH set.

General Discussion

Our results bear on two enduring issues in the word recognition literature: How lexical knowledge influences PH processing, and the source of word-frequency effects. On the former issue, the fact that RTs among PHs that contain most of the letters that make up their BWs were slower than RTs to PHs with more distinct orthographies suggests that high-letter-overlap PHs (e.g., WHELT), activate lexical structures more strongly than lower overlap PH's, and the stronger activation *interferes* with classifying the stimulus as a nonword.

Given this interpretation, computational models of all stripes would predict that PHs that share most of their orthography with their BWs *should be sensitive to BWF*. Yet, they are not. To further investigate this rather extraordinary result, we identified 113 PHs from the Armstrong and Plaut (2016) PH set whose length is identical to their BW and differ from it by only one letter. Despite the very high orthographic similarity between these PHs and their BWs, the zero-order correlation between their “No” RTs and their BWFs was nonexistent ($r = -0.0048$, $b = -0.098$, $t = -0.05$, $p = .96$). This raises an obvious question: How frequency sensitive would lexical decision RTs be to the BWs themselves? The answer can’t be obtained from the Armstrong and Plaut (2016) study, as the BWs weren’t included in it, but we were able to obtain mean RTs for all 113 BWs from the English Lexicon Project database of lexical decision RTs (Balota et. al., 2007). The correlation between BW “Yes” RTs and BWF turned out to be an impressively large 0.578 ($b = -24.87$, $t = 28.97$, $p < .001$). If lexical representations are inherently frequency-sensitive, how is it that making contact with what are presumably the same representations produces no effect of frequency on PH “No” RTs, but a very strong effect on BW “Yes” RTs? We are currently exploring this puzzle through additional data collection and modeling efforts.

The second notable result from our analyses is that BWF effects emerged among PHs that are orthographically distinct from their BWs, such that RTs were faster for PHs derived from HF BWs than for PHs derived from LF BWs. This result is also difficult to reconcile with extant computational models. According to the models, PHs derived from high-frequency BWs generate stronger levels of activation in structures that represent lexical knowledge than PHs derived from low-frequency BWs. All other things being equal, RTs to PHs derived from high-frequency BWs should be slower than RTs to PHs derived from low-frequency BWs, not faster.

While the BWF effects we found seem at odds with the models, the effects replicate several previous findings with humans (Van Orden, 1991; Van Orden et al., 1992; Ziegler, Jacobs, & Kluppel, 2001; see also Jared & Seidenberg, 1991, for similar results in a binary semantic classification task). To account for their findings, Van Orden, Ziegler, and their colleagues proposed that when confronted with a PH, people retrieve the BW spelling from their repository of whole-word spelling knowledge and compare the BW spelling with the PH spelling (see also Besner & Davelaar, 1983). The discrepancy between the two spellings provides evidence that the PH is a nonword. Assuming further that lexical memory for the spelling of high-frequency BWs is stronger than for the spelling of low-frequency BWs (Abrams & White, 2010), the spell check is completed more quickly for PHs derived from HF BWs than those derived from LF BWs.

If PH processing features a spell check, we would expect it to be performed on all PHs, regardless of how many

letters they share with their BW. Why, then, did the spell check not produce a BW frequency effect among high-overlap PHs? We speculate that when most or all of letters making up the BW are present in the PH, the need to retrieve the correct spelling for the BW is reduced or eliminated, as the PH “primes” the BW spelling directly. Direct priming of the BW spelling via PH orthography removes the temporal penalty that otherwise accompanies the retrieval of the BW spelling pattern from lexical memory. Hence, the more BW letters are present in the PH, the less sensitive the spell-check procedure is to BWF.

To our knowledge, a spell check has yet to be implemented in a localist computational model. However, Harm and Seidenberg (2004) incorporated a spell-check mechanism into their triangle PDP model in the form of a backpropagation channel between the semantic and orthographic processing networks. The channel yields a pattern of activation across the units of the orthographic network that captures the network’s knowledge of BW spelling. In principle, the BW pattern could be compared to the actual bottom-up pattern generated by the PH itself, and discrepancies between the two could be used to distinguish the PH spelling from the BW spelling. Through simulation, Harm and Seidenberg found that the higher the PH’s BWF, the more strongly and accurately the orthographic activation pattern produced by the backpropagation channel resembled the activation pattern generated by the BW itself. In other words, the backpropagation channel yielded the “raw material” – better “memory” for the spelling of a HF compared to a LF word – that could account for the sensitivity of the spell check to BWF.

Whether the BWF sensitivity of Harm and Seidenberg’s (2004) spell check would be attenuated among PHs that share most of their spelling with their BW can only be answered definitely through an actual simulation. However, we note that Harm and Seidenberg’s simulations also established that PHs that differ from their BW by only a single letter generate semantic activation patterns that more broadly match the patterns generated by their BWs than the patterns generated by more orthographically distinct PHs. Assuming that Harm and Seidenberg’s backpropagation channel would pass this difference on to the orthographic network, the orthographic activation pattern would more strongly and accurately match the BW orthographic pattern for high-overlap PHs than for lower-overlap PHs. In previous work, Plaut, McClelland, Seidenberg, and Patterson (1996) showed that if two activation “drivers” (such as orthographic overlap and BWF) jointly influence the rate at which the strength of an activation pattern grows over time, one driver can squash the influence of the other, yielding subadditive effects of the two on the growth function. Thus, previous PDP model results suggest that an expanded set of simulations could reveal a more frequency-sensitive spell check for low-overlap PHs than for high overlap PHs.

Acknowledgments

This work was supported by NSERC Grants RGPIN-2017-06310 to Blair Armstrong and RGPIN-2020-05506 to Michael Reynolds, and Grant AO998 to Derek Besner. Robert McCann would like to acknowledge the financial support of the United States Office of Personnel Management and Social Security Administration.

References

- Abrams, L., & White, K. K. (2010). Influences of word frequency, context, and age on spelling. In B. C. Fabini (Ed.), *Spelling skills: Acquisition, abilities and reading connection*. Hauppauge, NY: Nova Science.
- Armstrong, B.C., & Plaut, D.C. (2016). Disparate semantic ambiguity effects from semantic processing dynamics rather than qualitative task differences. *Language, Cognition and Neuroscience*, 31, 940-996.
- Balota, D., Yap, M. J., Cortese, M. J., Hutchison, K. A., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., & Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39, 445-459.
- Besner, D., & Davelaar, E. (1983). Suedohomophony effects in visual word recognition: Evidence for phonological processing. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 37, 300-305.
- Coltheart, M., Davelaar, E., Jonasson, T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and Performance VI*. New York: Academic Press.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). A dual-route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204-256.
- Forster, K. I. (2000). The potential for experimenter bias effects in word recognition experiments. *Memory & Cognition*, 28, 1109-1115.
- Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, 111, 662-720.
- Jared, D., & Seidenberg, M. (1991). Does word identification proceed from spelling to sound to meaning? *Journal of Experimental Psychology: General*, 120, 358-394.
- Kucera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, R.I.: Brown University Press.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10, 707-710.
- Lund, K., & Burgess, C. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods*, 28, 203-208.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- McCann, R.S., Besner, D., & Davelaar, E. (1988). Word recognition and identification: Do word-frequency effects reflect lexical access? *Journal of Experimental Psychology: Human Perception and Performance*, 14, 693-706.
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165-178.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56-115.
- Rodd, J., Gaskell, G., & Marslen-Wilson, W. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, 46(2), 245-266.
- Van Orden, G. C. (1987). A ROWS is a ROSE: Spelling, sound, and reading. *Memory & Cognition*, 15, 181-198.
- Van Orden, G. C. (1991). Phonological mediation is fundamental to reading. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition*. Hillsdale, NJ: Erlbaum.
- Van Orden, G. C., Stone, G. O., Garlington, K. L., Markson, L. R., Pinnt, G. S., Simonfy, C. M., & Bricchetto, T. (1992). "Assembled" phonology and reading: A case study in how theoretical perspective shapes empirical investigation. In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology, and meaning*. Amsterdam: North Holland.

- Wagenmakers, E. J., Steyvers, M., Raaijmakers, J. G., Shiffrin, R. M., Van Rijn, H., & Zeelenberg, R. (2004). A model for evidence accumulation in the lexical decision task. *Cognitive Psychology*, 48, 332-367.
- Weber, R. M. (1970). A linguistic analysis of first-grade reading errors. *Reading Research Quarterly*, 5, 427-451.
- Yarkoni, T., Balota, D., & Yap, M. (2008). Moving beyond Coltheart's *N*: A new measure of orthographic similarity. *Psychonomic Bulletin & Review*, 15, 971-979.
- Ziegler, J. C., Jacobs, A. M., and Kluppel, D. (2001). Pseudohomophone effects in lexical decision: Still a challenge for current word recognition models. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 547-559.