Reading Speech and Hearing Print: Constraining Models of Visual Word Recognition by Exploring Connections With Speech Perception

RON BOROWSKY, WILLIAM J. OWEN, and NIKKI FONOS, *University of Saskatchewan*

Abstract Current models of reading and speech perception differ widely in their assumptions regarding the interaction of orthographic and phonological information during language perception. The present experiments examined this interaction through a two-alternative, forced-choice paradigm, and explored the nature of the connections between graphemic and phonemic processing subsystems. Experiments 1 and 2 demonstrated a facilitation-dominant influence (i.e., benefits exceed costs) of graphemic contexts on phoneme discrimination, which is interpreted as a sensitivity effect. Experiments 3 and 4 demonstrated a symmetrical influence (i.e., benefits equal costs) of phonemic contexts on grapheme discrimination, which can be interpreted as either a bias effect, or an equally facilitative/inhibitory sensitivity effect. General implications for the functional architecture of language processing models are discussed, as well as specific implications for models of visual word recognition and speech perception.

> And so to completely analyse what we do when we read would almost be the acme of a psychologist's achievements, for it would be to describe very many of the most intricate workings of the human mind, as well as to unravel the tangled story of the most remarkable specific performance that civilisation has learned in all its history (Edmund Burke Huey, 1908, p. 6).

Echoing a sentiment that can be dated at least as far back as Galileo (see Chomsky, 1997), Huey's words still resonate with researchers who study basic reading processes today. In the current research on basic reading processes, one of the most "tangled" issues is whether the computation of word pronunciations requires a *single* processing route, or *dual* processing routes, in order to effectively map spelling representations onto sound representations. In order to aid our discussion of this issue, we will be referring to a framework for basic reading and speech perception processes that represents only the essential subsystems required for distinguishing between the two competing classes of models (see Figure 1). Specifically, there are separate subsystems for representing three types of linguistic knowledge: (a) *orthographic* (i.e., spelling), (b) *phonological* (i.e., speech sound), and (c) *semantic* (i.e., conceptual/meaning).

Traditional "dual-route" models of reading (e.g., Coltheart, Curtis, Atkins, & Haller, 1993; Coltheart, Davelaar, Jonasson, & Besner, 1977; Paap, McDonald, Schvaneveldt, & Noel, 1987) postulate that there must be two nonsemantic routes to map orthography onto phonology, or in other words, two distinct ways to compute pronunciations from print. These models typically involve localist representations for lexical (and sublexical) knowledge, whereby a single unit will correspond to a known word (or a known grapheme or phoneme). One route can be thought of as a sight vocabulary (SV) route, where orthographic lexical (i.e., individual whole-word) representations are mapped directly onto phonological lexical representations (see Figure 1). This route computes pronunciations for familiar words with *consistent* spelling-sound correspondences (e.g., *won*), and is relied on to compute correct pronunciations for familiar words with *inconsistent* spelling-sound correspondences (e.g., *one*). The second route can be thought of as a phonetic decoding (PD) route, where sublexical graphemic representations (i.e., of a letter or group of letters) are first mapped onto the phonemic (i.e., basic speech sound) representations that they are most consistently paired with, before a naming response is generated. This route is only capable of computing correct pronunciations for consistent words, and it is considered to be relied on for naming novel words or nonwords (e.g., *wum*).

Current "single-route" models of reading (e.g., Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989; see also Masson & Borowsky, 1998) postulate that only one nonsemantic route is necessary to map orthography onto phonology (see Figure 1). These models involve distributed, highly interconnected ("connectionist") representations that represent lexical and sublexical knowledge in weights associated with links that connect a set of processing units to one another, and

Canadian Journal of Experimental Psychology, 1999, 53:4, 294-305

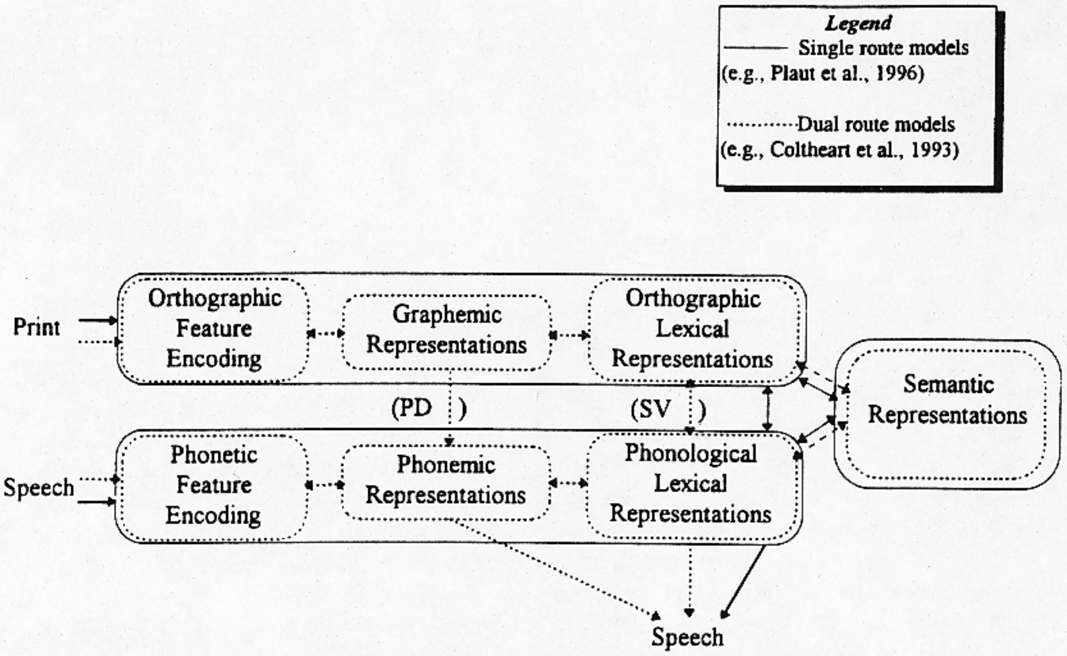

Figure 1. A framework for studying basic reading and speech perception processes.

instantiate this knowledge by evoking a unique pattern of activation across the processing units. Like the dual-route models, these models have been developed to capture diverse reading behaviours (e.g., the ability to name words with consistent or inconsistent spelling-sound correspondences, task dissociations with semantic ambiguity effects, nonword naming), but without requiring separate lexical representations for each word.

SINGLE VERSUS DUAL ROUTES, OR CONNECTIONIST-DISTRIBUTED VERSUS RULE-GOVERNED LOCALIST REPRESENTATIONS?

Proponents of single-route and dual-route models have typically treated this difference in number of nonsemantic routes as a necessary consequence of whether one subscribes to connectionist-distributed representations (e.g., Plaut et al., 1996) or rule-governed localist representations (e.g., grapheme-phoneme correspondence rules in Coltheart et al.'s 1993 model), making tests between these two types of models a rather complicated enterprise. However, the distinction between rule-governed localist and connectionist-distributed representations need not be a defining difference. Indeed, one could argue that the localist/distributed-connectionist distinction refers to a macro/micro distinction between two levels of description of the same psychological process (i.e., macro as the *computational* level and micro as either the *implementational* level or *representational and algorithmic* level; see Broadbent, 1985; Marr, 1982, and Rumelhart & McClelland, 1985, for similar arguments). Consider Zorzi, Houghton, and Butterworth's (1998) recent hybrid model of reading that clearly illustrates the advantages of implementing *connectionist* representations (which happen to be distributed representations in this model) within a dual-route architecture, accounting for a diverse set of normal, dyslexic, and developmental reading behaviours *without* any explicit grapheme-phoneme correspondence rules. This model demonstrates that the type of representation need not determine the number of processing routes between orthographic and phonological representation subsystems. Thus, the defining difference between single-route and dual-route models is, as their names imply, the number of routes postulated for effectively mapping orthography onto phonology. If the macro/micro distinction between levels of description is appropriate, then issues of representation may have been somewhat of a red herring that was confounded with single- and dual-route models.

NUMBER OF ROUTES AS A CONSEQUENCE OF GROUPING SUBSYSTEMS TOGETHER

We prefer to conceptualize the defining difference in the number of routes between the single-route and dual-route models as whether different subsemantic (i.e., preconceptual) levels of representation are grouped together or not. Figure 1 illustrates how the single-route models typically group the orthographic feature encoding level, the graphemic level, and the orthographic lexical level together into a single orthographic system, and, similarly, group the phonetic feature encoding level (in spirit, if not explicitly), the phonemic level, and the phonological lexical level into a single phonological system (as an exception, however, see Seidenberg, Petersen, MacDonald, & Plaut, 1996, for an example of a phonological subsystem that has "ungrouped" articulatory units, but see also Borowsky & Masson, 1999, for experiments that compromise Seidenberg et al.'s evidence for this modification to their model). With such grouping or "integration" (Plaut et al., 1996), it follows that there can be only one processing route that maps orthography onto phonology.

In contrast, the "ungrouped" dual-route models subscribe to a sublexical PD route (e.g., the mapping of graphemes onto phonemes in the Coltheart et al. 1993 model), and a separate lexical SV route (e.g., the mapping of whole-word orthography onto whole-word phonology in the Coltheart et al. model). Although much evidence has been reported in support of some aspects of the dual-route architecture (e.g., independent, but not exclusive, PD and SV routes, supported by the presence of developmental and acquired cases of pure and mixed forms of surface and phonological dyslexia, see Castles & Coltheart, 1993; Borowsky, McDougall, MacKinnon, & Hymel, 1999; McCarthy & Warrington, 1986; inadequate nonword naming performance of single-route simulation models, see Besner, Twilley, McCann, & Seergobin, 1990), some of this evidence has been scrutinized and accounted for by proponents of single-route models (e.g., Manis, Seidenberg, Doi, McBride-Chang, & Petersen, 1996; Plaut et al., 1996; but see Zorzi et al., 1998, for criticisms of some of these accounts). Clearly, this ongoing debate between proponents of single-route and dual-route models has provided a vigorous environment to test and develop both classes of models. In what follows, we argue that conceptualizing the difference between these models in terms of how subsemantic processing systems are grouped and connected together will provide a new perspective to model testing and development.

The framework in Figure 1 illustrates this "subsystem grouping" distinction between single-route and dual-route models of visual word recognition, as well as the importance of studying reading processes in the context of a larger language-processing framework. For example, the route that maps graphemic representations onto phonemic representations in dual-route models clearly serves to engage some of the subsystems that are involved in processing speech. Unfortunately, research on basic reading processes has been so focused on reading that the interconnections with other language processes are often dealt with only in a very restricted way (but see Kay, Lesser, & Coltheart, 1996, for a broader, applied language assessment model). The most basic defining difference between the single-route and dual-route architectures, the existence of a dedicated, feedforward processing connection that maps sublexical graphemic representations onto sublexical phonemic representations, necessarily involves representations of basic speech sounds. The research described in this paper involves a paradigm that will serve to inform models of reading with respect to such connections, within a context of a larger language processing framework that will also serve as a basis for studying how reading interacts with other language-processes, such as speech perception.

A PARADIGM FOR STUDYING INTERACTIONS BETWEEN LANGUAGE-PROCESSING SYSTEMS

The present research examines the nature and direction of interactions between sublexical processing systems involved in reading and speech perception, with a long term objective of examining how these systems interact with processing systems that are responsible for face and object recognition (i.e., mouth movements during speech production, e.g., Massaro & Stork, 1998; McGurk & MacDonald, 1976; pictures of objects, e.g., Masson & Borowsky, 1998). Masson and Borowsky have presented some of the background work for this program of research by examining the utility of signal detection measures of sensitivity for the purpose of studying the nature and direction of interactions between higher-level semantic processes and lower-level orthographic processes (see Fodor, 1983, for discussion on other modularity issues related to the implications of top-down processing). Using a combination of empirical and simulation work, Masson and Borowsky concluded that signal detection sensitivity can not be assumed to strictly represent sensitivity at a particular *sublexical* level when one is using stimuli that can be discriminated at a *lexical* (or higher, semantic) level of processing (despite arguments by previous researchers, e.g., Farah, 1989). As a result, the present research focuses on sublexical processing and precludes higher-level processes through the use of sublexical stimuli (e.g., the visually presented grapheme *ta* and auditory presentation of the speech syllable /ta/) and sublexical discrimination tasks (e.g., phonemic discriminations such as "did you hear /ta/ or /da/?", and graphemic discriminations such as "did you see *ta* or *da*?").[1]

The present research utilizes a two-alternative, forced-choice (2AFC) paradigm for examining sensitivity effects (inspired by Ratcliff & McKoon, 1997). The experiments reported here involve presenting a "context" stimulus (e.g., *ta*) simultaneously with a target in a different modality that is either congruent (e.g., /ta/), incongruent, or irrelevant to the context (i.e., a baseline), where the distinction between incongruent and irrelevant is determined by the two-alternative probe presented to the subject (e.g., context: *da*, target: /ta/, probe for incongruent: *heard* /ta/ or *heard* /da/, probe for irrelevant: *heard* /ta/ or *heard* /na/).

In this paradigm, a distinction is made between the presence of a bias effect, and the presence of a sensitivity effect over and above any effect of bias (Ratcliff & McKoon, 1997). A bias effect (e.g., an influence of the context on the participant's *interpretation* of what they perceived, or their

---

[1] For her honours thesis, Nikki Fonos examined the influence of graphemic contexts on phoneme perception using a signal detection paradigm, and obtained significant effects of context (congruent, incongruent) on both sensitivity (A′) and bias (B''_D). We subsequently found a better paradigm for evaluating facilitation versus inhibition than was possible with the signal detection approach, so only the more recent experiments are reported in the present paper.

willingness to choose one of the probe alternatives) is identified by an effect of the context on 2AFC accuracy that is *symmetrical* with respect to the irrelevant/baseline condition. Thus, a difference in accuracy between the congruent and baseline condition that is equal to the difference in accuracy between the baseline and incongruent condition would simply constitute a bias effect — that is, the participant's selection between the probe stimuli is influenced by the context stimulus if it corresponds to one of the probe alternatives. We note that Ratcliff and McKoon's supposition, whereby a symmetrical effect should be interpreted as a bias effect, can be challenged. For example, equally weighted facilitative and inhibitory connections *from* the context modality subsystem *to* the target modality subsystem, or from graphemic processing to phonemic processing in the present example, would also be consistent with the symmetrical effect. For present purposes, it suffices to say that such a symmetrical effect does not distinguish between a simple bias account and equally facilitative/inhibitory connections.

More importantly, the presence of a sensitivity effect (e.g., a direct influence of the context modality subsystem on the rate of information uptake in the target modality subsystem, or alternatively, on the baseline activation or threshold level in the target modality subsystem) can be identified in this paradigm as any significant deviation from the symmetrical effect of context on 2AFC accuracy (Ratcliff & McKoon, 1997). We have extended this logic to address facilitation-dominant versus inhibition-dominant cross-modality sensitivity effects (Borowsky, Fonos, & Owen, 1998; see also Becker, 1980). For example, if the benefits of a congruent context on 2AFC accuracy significantly exceed the costs of an incongruent context, a plausible account would be that there are facilitation-dominant connections from the context modality subsystem to the target modality subsystem. In other words, evidence for facilitation-dominant connections implies that the facilitative connections must outweigh (i.e., carry more influence than) the inhibitory connections (i.e., requiring a modification to interactive-activation architectures, e.g., McClelland & Rumelhart, 1981, that subscribe to equally weighted facilitative and inhibitory connections). Alternatively, if the costs of the incongruent condition on 2AFC accuracy significantly exceed the benefits of the congruent condition, a plausible account would be that there are inhibition-dominant connections from the context modality subsystem to the target modality subsystem. These inhibition-dominant connections could be implemented in interactive-activation architectures as inhibitory connections that outweigh the facilitative ones.

Instead of one type of connection "outweighing" the other, one might speculate whether evidence for a particular type of dominance could imply that there only exists connections of the dominant type (e.g., only facilitative or only inhibitory connections). While it is possible to describe many basic language processes without recourse to inhibitory connections (i.e., at a macro level of description, e.g., Borowsky & Masson, 1996; Masson & Borowsky, 1998), precluding facilitative connections would seem to be functionally implausible (i.e., how would activation spread from a low-level subsystem towards a higher-level subsystem if there were no facilitative connections?). Furthermore, precluding either type might seem to be neurophysiologically implausible (i.e., at a micro level of description) given that we have known about the presence of both facilitative and inhibitory synaptic connections for quite some time (e.g., Eccles, 1964; Sherrington, 1906), as well as their relevance to the receptive fields of simple and complex cortical cells in area V1 for basic visual feature encoding (e.g., Hubel, 1982; Hubel & Wiesel, 1959). However, we note that our earlier discussion regarding levels of description applies here as well, and thus "connections" can be described at macro and/or micro levels, and the existence of inhibitory connections at a micro level need not require inhibitory connections at the macro level of description.

WHAT CROSS-MODAL EFFECTS CAN CURRENT MODELS OF BASIC READING AND SPEECH PERCEPTION ACCOMMODATE?

Many models of basic reading processes have implemented theoretical *assumptions* on the nature and direction of interactions between processing subsystems. As mentioned above, dual-route models of reading advocate that there exists a *feedforward*, dedicated processing connection that maps sublexical graphemic representations onto sublexical phonemic representations, either in a way that is facilitative only (e.g., Coltheart et al., 1993), or both facilitative and inhibitory (e.g., Zorzi et al., 1998), but without any *feedback* connection. In contrast, most single-route models tend to adhere to *bidirectional* connections between grouped processing subsystems (e.g., fully recurrent connections between the orthographic subsystem and the phonological subsystem, Plaut et al., 1996; Seidenberg & McClelland, 1989; but see Masson & Borowsky, 1998, for an exception). Thus, dual-route models would easily account for a facilitation-dominant influence (i.e., a sensitivity effect) of grapheme contexts on phonemic discriminations (reflecting the degree to which this route has been established from reading experience) accompanied by a bias effect of phoneme contexts on graphemic discriminations, whereas fully recurrent single-route models that adhere to bidirectional connections between processing systems would, *in their current form*, most easily account for an equal, bidirectional influence of grapheme contexts on phoneme discriminations and vice versa (be it facilitation-dominant, inhibition-dominant, or equally facilitative and inhibitory). However, to be fair to these models of visual word recognition, many of them have focused strictly on visual word recognition and have not explicitly addressed the influence of speech perception processes on basic reading processes. As such, it would be premature to claim that a particular model could *never* account for a particular pattern of results. Nonetheless, as Zorzi et al. (1998) have pointed out, theoretical assumptions about the nature and direction of connections between processing subsystems have often been made on the basis of "instinct" or "trial and error" (see also McCloskey, 1991). It is essential to empirically constrain our models of visual word recognition with respect to the nature and direction of connections to and from other language subsystems.

Some models of speech perception have branched out so as to incorporate basic reading processes, and these models do make testable predictions regarding cross-modal influences. For example, in Fowler and Dekle's (1991) Direct-Realist Theory, which is based on Liberman and Mattingly's (1985) Motor Theory, phoneme perception is based on knowledge about the different vocal tract gestures that produce phonemes. This theory proposes that speech perception can only be affected by another source of information if it is considered (by the observer) to be emanating from the same speech source. Accordingly, there should be no influence whatsoever of grapheme contexts on phoneme detection, or vice versa. Indeed, Fowler and Dekle failed to reject the null hypothesis that the categorization of synthesized phonemes is unaffected by concurrent processing of congruent and incongruent graphemes (i.e., $p = .06$). However, it is important to note that their results were consistent with the existence of a small effect.

In Massaro, Cohen, and Thompson's (1988) Fuzzy Logical Model of Perception, phonemes are initially processed independently, then integrated with graphemes during the matching of features to prototypes in memory. Phoneme perception is accomplished by selecting the syllable prototype that best matches the collection of features. In support of their theory, Massaro et al. reported a small, but significant, effect of concurrent grapheme contexts on the categorization of phonemes. However, Massaro et al.'s model clearly predicts that the cross-modal influences of graphemic and phonemic processing should be restricted to bias effects, and not sensitivity effects; these effects were not separated in their study.

Finally, Grainger and colleagues (Grainger & Ferrand, 1994; Jacobs, Rey, Ziegler, & Grainger, 1998) have developed an interactive-activation model that represents a merging of McClelland and Elman's (1986) TRACE model of speech perception with McClelland and Rumelhart's (1981) interactive-activation model of basic reading processes. The framework for this model thus includes both facilitative and inhibitory connections between orthographic and phonological representations at both the sublexical and lexical levels. In this model, phoneme and grapheme perception occurs when a criterial level of activation is reached in the sublexical units that represent a particular phoneme. Given this framework, an equal, bidirectional influence of grapheme contexts on phoneme discriminations and vice versa (be it facilitation-dominant, inhibition-dominant, or equally facilitative and inhibitory) could easily be accounted for (see Jacobs et al., 1998, and Ziegler, Van Orden, & Jacobs, 1997, for cogent research on bidirectional consistency effects).

The first two experiments in this paper examine the influence of graphemic contexts on phoneme discrimination; the remaining two experiments examine the influence of phonemic contexts on grapheme discrimination.

Experiments 1 and 2

Experiments 1 and 2 investigated the influence of a graphemic context upon phoneme detection. Experiment 1 was designed to be a relatively difficult phoneme discrimination task. In order to evaluate whether the pattern of results would change as a function of location on the accuracy scale (i.e., a scaling account), or as a function of target discriminability adding to bias (i.e., a target bias account), the discrimination was made easier in Experiment 2 by decreasing the degradation of the phoneme targets.

METHOD

Participants. Twenty-four University of Saskatchewan students participated for partial credit in an introductory psychology class in Experiment 1, and another 38 students from the same population participated in Experiment 2. All reported English as their first language and normal (or corrected-to-normal) vision and hearing.

Apparatus. An IBM-compatible computer with Micro Experimental Laboratory (MEL) software controlled the timing of the events and recording of the data. Graphemes were presented in white on a black background using a NEC colour monitor (Model JC-15W1VMA). A pair of Altec Lansing ACS5 speakers, placed on either side the monitor, were used to present the auditory stimuli via a Creative Lab Sound Blaster compatible 16-bit audio card. The "1" and "2" keys on the numeric key pad were used to collect participants' responses. Twenty-one participants in Experiment 2 used the same apparatus described above, while the remainder used a similar apparatus, with the only differences being that a different NEC colour monitor (Model JC-1433VMA) and AST (Advantage series) speakers were used.

Materials and design. Three phonemes (i.e., syllables), /ta/, /na/, and /da/, matched on place of articulation (i.e., alveolar) were used in this set of studies. Creative WaveStudio (version 2.0) was used to record the phonemes (spoken by a male). The three phonemes were edited such that each initial onset was added to the same /a/ phoneme. Each phoneme was recorded in 16-bit mono, at a sampling frequency of 22KHz, and was 500 ms in duration. Each phoneme stimulus was presented simultaneously with white-noise (the MEL white-noise generator was set to 88% maximum output in Experiment 1, and reduced to 86% maximum output in Experiment 2). MEL code specification for generation of the white-noise output was AUDIO_SET_VOLUME(4, 0, 88) and AUDIO_SET_VOLUME(4, 0, 86) for Experiments 1 and 2, respectively. The sound pressure level of the auditory stimuli was measured by a Radio Shack digital sound level meter (Model 33-2055) directly in front of one of the speakers. The maximum ambient sound pressure level in the testing room was 66 dB. The auditory stimuli with white-noise were presented at a maximum of 70 dB and 69 dB in Experiments 1 and 2, respectively. Without white-noise, the auditory stimuli produced a maximum sound pressure level of 68 dB.

Three presentation conditions were created based on the match of the grapheme context to the phoneme target and the response probes. The grapheme context stimulus was presented simultaneously with the phoneme target stimulus and was congruent, incongruent, or irrelevant to the target. In the *congruent* condition, the grapheme context matched the phoneme target, and the visually presented response probe for this condition contained the target and one of the other two stimuli (e.g., presentation of the grapheme context *ta* and phoneme target /ta/, probed with *heard ta or heard da*). In the *incongruent* condition, the grapheme context and phoneme target did not match, and the response probe contained both the context and target stimuli (e.g., presentation of the grapheme context *da* and phoneme target /ta/, probed with *heard ta or heard da*). In the *irrelevant* condition, the grapheme context and phoneme target did not match, and the response probe contained the target and the irrelevant remaining stimulus (i.e., not the context grapheme; e.g., presentation of the grapheme context *na* and phoneme target /ta/, probed with *heard ta or heard da*). The three phonemes and graphemes appeared in each of the congruent, incongruent, and irrelevant conditions equally often, and the correct alternative of the response probe appeared equally often on the right- or left-hand side, creating 36 trial conditions. The experiment consisted of 11 continuous blocks of the 36 randomized trial conditions, with the first block of 36 trials serving as practice trials.

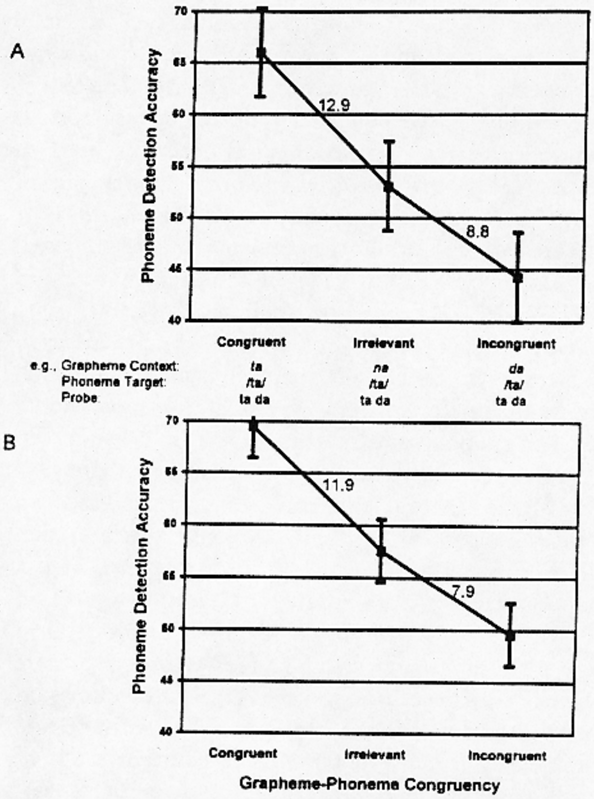

Figure 2. Mean phoneme detection accuracy (in percent) as a function of grapheme-phoneme congruency for: (A) Experiment 1, and (B) Experiment 2. Confidence intervals were calculated using the formula for a within-subjects design as outlined in Loftus and Masson (1994).

Procedure. Participants were tested individually in a quiet laboratory. They were instructed, both verbally and in writing, that they would see a grapheme (e.g., *ta*, *na*, or *da*) in the middle of the computer screen and, at the same time, they would hear a phoneme presented in white-noise. They were told to pay attention to both what they saw and what they heard (and that sometimes they would match, sometimes not), but to respond to what they heard, selecting from a two-alternative response probe, as quickly and accurately as they could. If the participant was unsure of what they had heard, they were told to guess. The sequence of events was: (1) a fixation mark appeared in the centre of the screen, (2) the participant pressed the space-bar to initiate each trial, (3) after a 100-ms interstimulus interval (ISI) a clearly visible grapheme appeared in the centre of the screen simultaneously with the degraded phoneme target speech syllable, both for a total of 500 ms, (4) after a 100-ms ISI, a two-alternative response probe was presented visually, in bright text, near the bottom of the screen (e.g., *heard ta [press1]*, *heard da [press 2]*), and (5) accuracy was coded immediately by the computer. The procedure was approximately 20 minutes in duration, during which time the experimenter remained in the laboratory.

RESULTS

Experiment 1. Overall mean response accuracy for the congruent, irrelevant, and incongruent conditions is presented in Figure 2A. Repeated measures analyses of variance (ANOVA) of condition (congruent, irrelevant, and incongruent) on accuracy was significant, $F(2,46) = 26.69$, $MSE = 2{,}858.71$, $p < .001$.

Dependent t-tests showed that mean accuracy for the congruent condition was significantly greater than that for the irrelevant condition, $t(23) = 5.72$, $SE = 2.25$, $p < .001$, and the irrelevant condition mean accuracy was significantly greater than the incongruent condition mean accuracy, $t(23) = 3.92$, $SE = 2.25$, $p < .01$. The test of the difference between response accuracy in the congruent condition minus the irrelevant condition (12.9%) and the irrelevant condition minus the incongruent condition (8.8%) was also significant, $t(23) = 2.15$, $SE = 1.89$, $p < .05$.

Experiment 2. No differences in accuracy were found between the two groups of participants run on the two sets of apparatus (all $ts < 1.38$, $ps > .17$) and, therefore, the two data sets were combined. Overall mean response accuracy for the congruent, irrelevant, and incongruent conditions is presented in Figure 2B.

A repeated measures ANOVA of condition (i.e., congruent, irrelevant, and incongruent) on accuracy in Experiment 2 was significant, $F(2,74) = 46.07$, $MSE = 3{,}775.26$, $p < .001$. Dependent samples t-tests indicated that response accuracy in the congruent condition was significantly greater than in the irrelevant condition, and that the irrelevant condition accuracy was significantly greater than for the incongruent condition (all $ts > 4.90$, $ps < .001$). The difference of the congruent condition minus the irrelevant condition accuracy (11.9%) and the irrelevant condition minus the incongruent condition accuracy (7.9%) was significant, $t(37) = 2.23$, $SE = 1.78$, $p < .05$.

The purpose of Experiment 2 was to determine if an increase in response accuracy would alter the facilitation-dominant effect found in Experiment 1. Specifically, a scaling account (whereby the incongruent condition is influenced by a floor effect that limits below-chance performance) would be plausible if the facilitation-dominant effect changed to either a symmetrical effect or an inhibition-dominant effect when accuracy is higher. Alternatively, a target bias account (as suggested by a reviewer, whereby the target itself contributes to additional bias in the congruent condition) would be plausible if the facilitation-dominant effect became more pronounced when accuracy is higher. An independent samples t-test was conducted to confirm that the response accuracy for the baseline (i.e., irrelevant) condition in Experiment 2 was significantly greater than that observed in Experiment 1, $t(60) = 2.80$, $SE = 1.59$, $p < .01$. To determine if the accuracy pattern of results differed between Experiments 1 and 2, an ANOVA of condition by experiment was conducted on the accuracy data. There was a significant main effect of experiment, $F(1,60) = 12.83$, $MSE = 871.69$, $p < .01$, and of condition, $F(2,120) = 70.03$, $MSE = 6{,}413.74$, $p < .001$. There was no interaction between experiment and condition ($F < 1.00$), and thus Experiment 1 and 2 accuracy data were combined. A one-sample t-test comparing the difference of the facilitation effect (i.e., congruent minus irrelevant condition accuracy) minus the inhibition effect (i.e., irrelevant minus incongruent condition accuracy) to a mean of zero was conducted. In this analysis, all negative numbers were replaced with a zero value so as to not artificially inflate any differences; however, the untransformed data yielded the same pattern of results. This difference score was significantly greater than zero, $t(61) = 3.28$, $SE = 1.12$, $p < .01$, and the 95% confidence intervals ranged from +1.42% to +5.89% (see Table 1).

TABLE 1
Mean Difference Between Facilitation and Inhibition (in percent), and 95% Confidence Intervals as a Function of Task

Task	Facilitation Minus Inhibition	95% Confidence Interval	
	Mean Effect	Lower Bound	Upper Bound
Phoneme Detection (Experiments 1 and 2)	+3.6	+1.42	+5.89
Grapheme Detection (Experiments 3 and 4)	- 1.1	- 3.60	+1.27

DISCUSSION

The significant facilitation-dominant influence of a concurrent grapheme context on phoneme target discrimination accuracy in Experiment 1 suggests that phoneme discrimination sensitivity benefits from an effect of congruent grapheme contexts that exceeds the cost incurred by incongruent graphemes. This result could be accommodated by connections from grapheme representations to phoneme representations that are primarily facilitative (i.e., facilitative connections that dominate over any inhibitory connections). Before accepting this facilitation-dominant pattern, however, Experiment 2 was conducted to check whether the effect is an artefact of a low overall level of accuracy in the experiment.

Despite a significant increase in accuracy in Experiment 2 over Experiment 1, the same facilitation-dominant influence of grapheme contexts on phoneme discrimination was still observed. This suggests that the facilitation-dominant influence of grapheme contexts on phoneme discrimination is not compromised by a scaling effect on overall accuracy, nor any form of additional bias from the target as a function of increased discriminability.

We now turn to experiments that explore the influence of phonemic contexts on grapheme discrimination.

Experiments 3 and 4

Experiments 3 and 4 investigated the influence of a phonemic context on grapheme detection. Experiment 3 was designed to be a relatively difficult grapheme discrimination task. In order to evaluate whether a scaling or target bias effect could be contributing to the pattern of results, the discrimination was made easier in Experiment 4 by decreasing the degradation of the grapheme targets.

METHOD

Participants. Twenty-four University of Saskatchewan

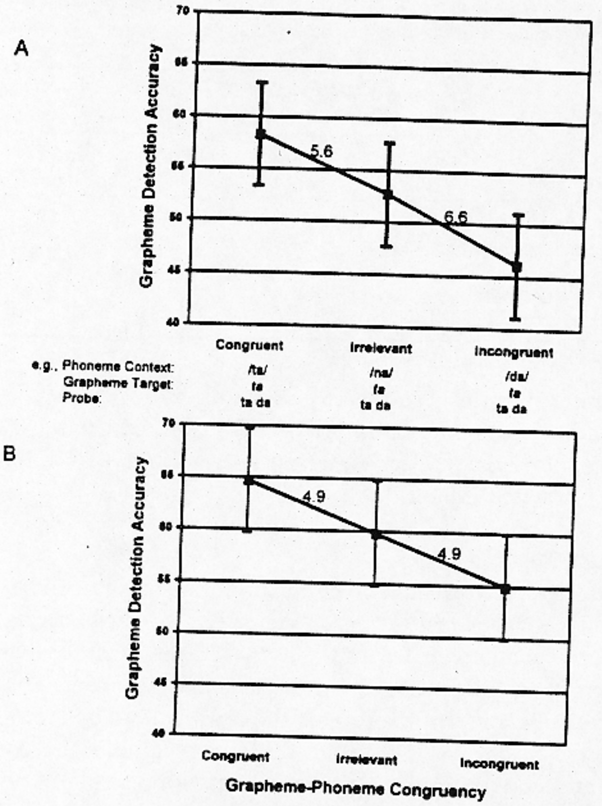

Figure 3. Mean grapheme detection accuracy (in percent) as a function of grapheme-phoneme congruency for: (A) Experiment 3, and (B) Experiment 4. Confidence intervals were calculated using the formula for a within-subjects design as outlined in Loftus and Masson (1994).

students were paid $5 for participating in Experiment 3, while another 24 students participated in Experiment 4 for partial credit in an introductory psychology class. All reported English as their first language and normal (or corrected-to-normal) vision and hearing.

Apparatus. The same equipment as reported in Experiment 1 was used. The two differences in this experiment were that the graphemes were presented in dark gray on a black background (to avoid ceiling effects on discrimination accuracy) and a MEL serial response box was used to initiate each trial and collect responses.

Materials and design. The same materials and design as in the previous experiments were used in this experiment. The only differences were that clearly audible phonemes (i.e., without any white-noise) now provided the context, and the graphemes were degraded by contrast-reduction and presented as targets. MEL code specification for the specific level of the contributions of red, green, and blue for dark gray was SET_PALETTE_VGA(8,5,5,6) and SET_PALETTE_VGA(8,6,6,6) in Experiments 3 and 4, respectively. The luminance for the background of the display, as measured by a Tektronic J6523-2, 1° narrow angle luminance probe attached to a J16 digital photometer at a distance of 40 cms from the computer screen, was .6 cd/m^2. The target graphemes were presented at a luminance of .7 cd/m^2 and .75 cd/m^2 for Experiments 3 and 4, respectively. Although this form of degradation is arguably different from the addition of noise used in Experiments 1 and 2, it has been shown to be suitable for demonstrating both facilitation and inhibition priming effects under appropriate baseline conditions in lexical decision (e.g., Borowsky & Besner, 1991, 1993). Furthermore, Borowsky and Besner (1993) showed that this form of degradation interacts in an overadditive fashion with priming in such a way that can not be attributed to a simple inhibition-based strategy (i.e., an expectancy-based account of this interaction was ruled out in favour of an activation-based account).

Procedure. The procedure was similar to that in Experiments 1 and 2 except that participants were to discriminate between target graphemes. In order to obtain a similar mean response accuracy for the baseline (i.e., irrelevant) condition in this grapheme discrimination task as observed in the same condition in the phoneme discrimination task in Experiment 1, a pilot study suggested that reducing the target presentation to 150 ms *and* degrading the graphemes were both necessary to increase the difficulty of discriminating the graphemes.

Participants were instructed, both verbally and in writing, that they would hear a clearly audible phoneme, and, at the same time, they would see a degraded grapheme in the middle of the computer screen. They were told to pay attention to both what they heard and what they saw (and that sometimes they would match, sometimes not), but to respond to what they saw, selecting from a two-alternative response probe, as quickly and accurately as they could. If the participant was unsure of what they had seen, they were told to guess. The sequence of events was: (1) a fixation mark appeared in the centre of the screen, (2) the participant pressed the middle key on the response box to initiate each trial, (3) after a 100-ms ISI, a degraded grapheme appeared in the centre of the computer screen for 150 ms during the simultaneous presentation of a phoneme for 500 ms (i.e., the grapheme and phoneme onsets were simultaneous), (4) after a 200-ms ISI, a two-alternative response probe was presented visually, in bright text, near the bottom of the screen (e.g., *saw ta [press1]*, *saw da [press 2]*) using the keys labelled "1" or "2" on the response box, and (5) accuracy was coded immediately by the computer. The experiment consisting of 396 trials with the first 36 trials considered as practice items.[2]

---

[2] We initially thought that we would have to individually set the level of visual degradation for each participant, thus, they were given 20 pre-test trials conforming to the experimental procedure described above. This was followed by 36 trials in which only a degraded grapheme was presented followed by a two-alternative response probe

The procedure took about 25 minutes, during which time the experimenter remained in the laboratory.

RESULTS

Experiment 3. The mean response accuracy data for the congruent, irrelevant, and incongruent conditions are presented in Figure 3A. A repeated measures ANOVA of condition on accuracy was significant, $F(2,46) = 5.58$, $MSE = 888.03$, $p < .05$. Dependent t-tests showed that mean accuracy in the congruent condition was significantly greater than in the irrelevant condition, $t(23) = 2.38$, $SE = 2.35$, $p < .05$, which in turn was significantly greater than the incongruent condition, $t(23) = 2.14$, $SE = 3.07$, $p < .05$. The test of the difference between mean accuracy for the congruent condition minus the irrelevant condition (5.6%) and for the irrelevant condition minus the incongruent condition (6.6%) was not significant, $t(23) = -0.43$, $SE = 2.24$, $p = .67$.

Experiment 4. Overall mean response accuracy for the congruent, irrelevant, and incongruent conditions is presented in Figure 3B. A repeated measures ANOVA of condition (i.e., congruent, irrelevant, and incongruent) on accuracy was significant, $F(2,46) = 8.14$, $MSE = 579.35$, $p < .001$. Dependent samples t-tests indicated that response accuracy in the congruent condition was significantly greater than in the irrelevant condition, and that the irrelevant condition accuracy was significantly greater than for the incongruent condition (all $ts > 2.32$, $ps < .05$). The difference of the congruent condition minus the irrelevant condition accuracy (4.9%) and the irrelevant condition minus the incongruent condition accuracy (4.9%) was not significant, $t(23) = .02$, $SE = 2.27$, $p = .99$.

Given that the purpose of Experiment 4 was to determine if an increase in response accuracy would alter the symmetrical effect found in Experiment 3 (see Experiment 2 results section), an independent samples t-test was conducted to confirm that the response accuracy for the baseline (i.e., irrelevant) condition in Experiment 4 was greater than that observed in Experiment 3, $t(46) = -3.06$, $SE = 2.31$, $p < .01$. To determine if the pattern of results differed between Experiments 3 and 4, an ANOVA of condition by experiment was conducted on the accuracy data. There was a significant main effect of experiment, $F(1,46) = 12.22$, $MSE = 1{,}962.98$, $p < .01$, and of condition, $F(2,92) = 12.60$, $MSE = 1{,}450.13$, $p < .001$. There was no interaction between experiment and condition ($F < 1$), and thus Experiment 3 and 4 accuracy data were combined. A one-sample t-test comparing the difference of the facilitation effect (i.e., congruent minus irrelevant condition accuracy) minus the inhibition effect (i.e., irrelevant minus incongruent condition accuracy) to a mean of zero was conducted and the confidence intervals determined. In this analysis, all negative numbers were replaced with a zero value; however, the untransformed scores yielded the same pattern of results. There was no significant difference from zero, $t(47) = -.95$, $SE = 1.20$, $p = .34$, and the 95% confidence intervals ranged from -3.60% to +1.27% (see Table 1). In comparison to the confidence interval reported for Experiments 1 and 2 combined, this confidence interval does not overlap with it, and thus the combined phoneme detection experiments yield a facilitation-dominant effect that differs significantly from the symmetrical effect obtained with the grapheme detection experiments. An ANOVA of condition by task (collapsing Experiments 1 and 2 together, and 3 and 4 together) supports this analysis, as it showed a main effect of condition, $F(2,216) = 67.39$, $MSE = 6{,}759.02$, $p < .001$, no main effect of task ($F < 1.00$), and, most importantly, a significant interaction of condition by experiment, $F(2,216) = 6.61$, $MSE = 661.74$, $p < .01$.

DISCUSSION

The symmetrical influence of a concurrent phoneme context on grapheme target discrimination accuracy in Experiment 3 suggests that phoneme discrimination benefits from an effect of congruent grapheme contexts that equals the cost incurred by incongruent graphemes. This symmetrical influence of a concurrent phoneme context on grapheme target discrimination accuracy remained despite a significant increase in baseline accuracy in Experiment 4, suggesting that this effect is not compromised by a scaling effect on overall accuracy, nor any form of additional target bias as a function of increased discriminability.

As described in the introduction, this result could be accommodated either by: (1) connections from phoneme representations to grapheme representations that are equally facilitative and inhibitory, or (2) a bias effect with no direct connections between these subsystems.

General Discussion

Experiments 1 and 2 provided evidence for facilitation-dominant connections from the subsystem that represents graphemes to the subsystem that represents phonemes. Given that this pattern was found at two different levels of accuracy, both a scaling account and a target bias account of the data can be ruled out. The scaling account would maintain that the incongruent condition is influenced by a floor effect that limits below-chance performance, but this account would have been plausible only if the facilitation-dominant effect changed (e.g., became a symmetrical effect or an inhibition-dominant effect) when accuracy is significantly higher. The target bias account would maintain that the target itself contributes to additional bias in the congru-

---

from which participants decided what they saw or thought they saw. Based upon the two pre-tests, most participants received the same (moderate) level of degradation for the experimental trials. Those participants ($n = 5$) who received clearer or more degraded presentations were not included in the analysis.

ent condition, but this account would have been plausible only if the facilitation dominant effect became more pronounced when the target is easier to discern (i.e., when accuracy is significantly higher). Having ruled out these alternative accounts, the influence of the facilitative connections from graphemic representations to phonemic representations must therefore outweigh the influence (if any) of inhibitory connections.

Experiment 3 provided evidence for a symmetrical effect of phoneme contexts on grapheme discrimination accuracy, which could be attributed to either a bias effect or equally facilitative and inhibitory connections from the system responsible for processing phonemes to the system responsible for processing graphemes. Experiment 4 also provided evidence for the same symmetrical effect of phoneme contexts on grapheme discrimination when baseline accuracy performance is higher, thus ruling out any scaling or target bias influences.

The present results suggest that models involving graphemic and phonemic processing that clearly predict either no influence of one modality on the other, or only bias effects, are incorrect (i.e., Fowler & Dekle, 1991; Massaro et al., 1988). Models that could be considered generally as bidirectional interactive-activation or fully recurrent architectures (e.g., Jacobs et al., 1998; Plaut et al., 1996) are probably better equipped, in their current form, to handle equivalent bidirectional effects, unlike the unidirectional facilitation-dominant effect reported here. It remains to be seen whether these latter models could be modified to handle the present results. Adjusting the relative influence of the facilitative connections from the system representing graphemes to the system representing phonemes to exceed that of the inhibitory connections, while maintaining equal influence of the facilitative and inhibitory connections from the system representing phonemes to the system representing graphemes, seems like one possible way to simulate a facilitation-dominant pattern.

The facilitation-dominant influence of graphemic processing on phonemic processing, over and above any bias effects involving these subsystems, can be handled most easily by current models which hypothesise feedforward connections from the system that represents graphemes to the system that represents phonemes, and no feedback connections. This state of affairs is best captured by Coltheart et al.'s (1993) dual-route model, and Masson and Borowsky's (1998) distributed single-route model. Zorzi et al.'s (1998) hybrid connectionist dual-route model would require a single modification whereby the influence of the facilitative connections outweigh the inhibitory connections.

The present results thus indicate the type of connections that must exist between subsystems that process graphemic and phonemic information, regardless of whether the architecture involves single or dual routes. It is important to note, however, that the present framework and paradigm can be used to test a unique prediction of single-route models, which follows from the experiments reported here. Current single-route models do not distinguish between sublexical and lexical subsystems (i.e., they are grouped together, see Figure 1). Given that this grouping requires the same connections to be involved in processing both sublexical stimuli (e.g., graphemes, phonemes) and lexical stimuli (i.e., words), single-route models must predict that the same pattern of results obtained in the present experiments with sublexical stimuli will also be obtained when using lexical stimuli (i.e., a facilitation-dominant influence of printed word contexts on spoken word discrimination, and a symmetrical influence of spoken word contexts on printed word discrimination). Our future research will explore this prediction. We hope that other researchers will also find the framework and paradigm outlined here to be useful for developing their models of basic language processing.

This research was supported by grants from the Social Science and Humanities Research Council of Canada, and the Natural Sciences and Engineering Research Council of Canada to R. Borowsky, and an NSERC post-graduate scholarship to W. J. Owen. We thank Michael Masson and Jim Cheesman for insightful discussions, and Pete Dixon for comments on the manuscript. Address correspondence to Ron Borowsky, Department of Psychology, University of Saskatchewan, 9 Campus Drive, Saskatoon, Saskatchewan S7N 5A5 (E-mail: ron.borowsky@usask.ca; http://www.usask.ca/psychology/).

References

Becker, C. A. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory & Cognition, 8*, 493-512.

Besner, D., Twilley, L., McCann, R. S., & Seergobin, K. (1990). On the association between connectionism and data: Are a few words necessary? *Psychological Review, 97*, 432-446.

Borowsky, R., Fonos, N., & Owen, W. J. (1998, June). *On the modularity of orthographic and phonological perception.* Paper presented at the 8th annual meeting of the Canadian Society of Brain, Behavior, and Cognitive Science, Ottawa, Canada.

Borowsky, R., & Masson, M. E. J. (1996). Semantic ambiguity effects in word identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 63-85.

Borowsky, R., & Masson, M. E. J. (1999). Frequency effects and lexical access: On the interpretation of null pseudohomophone base-word frequency effects. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 270-275.

Borowsky, R., McDougall, P., MacKinnon, G. E., & Hymel, S. (1999). *On the diagnosis of surface and phonological dyslexias: Measuring reliance on sight vocabulary and phonetic decoding during real word recognition.* Manuscript submitted for publication.

Broadbent, D. (1985). A question of levels: Comment on McClelland and Rumelhart. *Journal of Experimental Psychology: General, 114,* 189-192.

Castles, A., & Coltheart, M. (1993). Varieties of developmental dyslexia. *Cognition, 47,* 149-180.

Chomsky, N. (1997). *Language and mind.* Princeton University Lecture (transcript).

Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading: Dual route and parallel-distributed-processing approaches. *Psychological Review, 100,* 598-608.

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and performance VI* (pp. 535-555). London: Academic Press.

Eccles, J. C. (1964). *The physiology of synapses.* Berlin: Springer-Verlag.

Farah, M. (1989). Semantic and perceptual priming: How similar are the underlying mechanisms? *Journal of Experimental Psychology: Human Perception and Performance, 15,* 188-194

Fodor, J. A. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.

Fowler, C. A., & Dekle, D. J. (1991). Listening with eye and hand: Cross-modal contributions to speech perception. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 816-828.

Grainger, J., & Ferrand, L. (1994). Phonology and orthography in visual word recognition: Effects of masked homophone primes. *Journal of Memory and Language, 33,* 218-233.

Hubel, D. H. (1982). Exploration of the primary visual cortex, 1955-1978. *Nature, 299,* 515-524.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology, 148,* 574-591.

Huey, E. B. (1908). *The psychology and pedagogy of reading.* New York: The MacMillian Company.

Jacobs, A. M., Rey, A., Ziegler, J. C., & Grainger, J. (1998). MROM-P: An interactive activation, multiple read-out model of orthographic and phonological processes in visual word recognition. In J. Grainger & A. M. Jacobs (Eds.), *Localist connectionist approaches to human cognition* (pp. 147-188). Mahwah, NJ: Erlbaum.

Kay, J., Lesser, R., & Coltheart, M. (1996). Psycholinguistic assessments of language processing in aphasia (PALPA): An introduction. *Aphasiology, 10,* 159-180.

Liberman, A., & Mattingly, I. (1985). The motor theory of speech perception revised. *Cognition, 21,* 1-36.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subjects designs. *Psychonomic Bulletin & Review, 1,* 476-490.

Manis, F. R., Seidenberg, M. S., Doi, L. M., McBride-Chang, C., & Peterson, A. (1996). On the bases of two subtypes of developmental dyslexia. *Cognition, 58,* 157-195.

Marr, D. (1982). *Vision.* San Francisco, CA: Freeman.

Massaro, D., Cohen, M., & Thompson, L. (1988). Visual language in speech perception: Lipreading and reading. *Visual Language, 22,* 8-31.

Massaro, D. W., & Stork, D. G. (1998). Speech recognition and sensory integration. *American Scientist, 86,* 236-244.

Masson, M. E. J., & Borowsky, R. (1998). More than meets the eye: Context effects in word recognition. *Memory & Cognition, 26,* 1245-1269.

McCarthy, R., & Warrington, E. K. (1986). Phonological reading: Phenomena and paradoxes. *Cortex, 22,* 359-380.

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

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive model of context effects in letter perception: I. An account of basic findings. *Psychological Review, 88,* 375-407.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114,* 159-188.

McCloskey, M. (1991) Networks and theories: The place of connectionism in cognitive science. *Psychological Science, 2,* 387-395.

McGurk, H., & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature, 264,* 746-748.

Paap, K. R., McDonald, J. E., Schvaneveldt, R. W., Noel, R. W. (1987). Frequency and pronounceability in visually presented naming and lexical decision tasks. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 221-243). Hillsdale, NJ: Erlbaum.

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.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review, 104,* 319-343.

Rumelhart, D. E., & McClelland, J. L. (1985). Levels indeed! A response to Broadbent. *Journal of Experimental Psychology: General, 114,* 193-197.

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

Seidenberg, M. S., Peterson, A., MacDonald, M. C., & Plaut, D. C. (1996). Pseudohomophone effects and models of word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 48-62.

Sherrington, C. S. (1906). *The integrative action of the nervous system.* New York: Scriber's (2nd ed.). New Haven, CT: Yale University Press, 1947.

Tabachnick, B. G., & Fidell, L. S. (1996). *Using multivariate statistics.* New York: HarperCollins College Publishers.

Ziegler, J. C., Van Orden, G. C., & Jacobs, A. M. (1997). Phonology can help or hurt the perception of print. *Journal of Experimental Psychology: Human Perception and Perfor-*

mance, 23, 845-860.

Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1131-1161.

Sommaire

Les modèles actuels de perception de l'écriture et du discours diffèrent largement dans leurs hypothèses quant à l'interaction de l'information orthographique et phonologique pendant la perception du langage. Les expériences actuelles ont permis d'examiner cette interaction grâce à un paradigme à deux solutions et à choix forcé, et d'explorer la nature des connexions entre les sous-systèmes de traitement à graphèmes et à phonèmes. Les expériences 1 et 2 ont démontré une influence à facilitation dominante (c.-à-d. que les avantages dépassaient les coûts) parmi les contextes à graphèmes sur la discrimination phonémique, ce qui est interprété comme un effet de sensibilité. Les expériences 3 et 4 ont démontré une influence symétrique (c.-à-d. que les avantages égalent les coûts) des contextes phonémiques sur la discrimination à graphèmes, ce qui peut être interprété soit comme un effet de biais, soit comme un effet de sensibilité aussi facilitant qu'inhibant. On traite des conséquences générales des modèles de traitement linguistique sur l'architecture fonctionnelle, ainsi que des conséquences spécifiques que peuvent avoir, sur les modèles, la reconnaissance visuelle des mots et la perception du discours.