Developing theories of priming with an eye on function.

Jeffrey S. Bowers
Department of Experimental Psychology
University of Bristol
Bristol, BS8-1TN
j.bowers@bris.ac.uk

Sid Kouider
Laboratoire de Sciences Cognitives et Psycholinguistique,
EHESS / CNRS,
54 Boulevard Raspail,
75270 Paris Cedex 06,
France.
sidk@lscp.ehess.fr

Running head: functional considerations

Long-term priming has been one of the most extensively studied phenomena in cognitive psychology during the past 20 years, with literally hundreds of papers devoted to this topic (for reviews, see Bowers, 2000a; Roediger & McDermott, 1993; Tenpenny, 1995). For the most part, this research has been carried out within a memory framework in which priming is understood as a form of implicit memory that is compared and contrasted with various expressions of explicit memory, in particular, recall and recognition. Within this framework, one of the central debates has been whether implicit and explicit memories are supported by common or separate systems, but whatever case, it is agreed that priming is mediated by a memory system whose central function is, presumably, memory. As a consequence, the constraints used in developing theories of priming are largely the priming data themselves along with related memory phenomena, and theories are judged a success to the extent they account for these data.

Although this framework has been successful in generating a rich database, we think it mischaracterizes the underlying function of the system(s) that supports priming. On our view, priming is a behavioral manifestation of learning processes embedded with perceptual (and sometimes conceptual) systems whose main function is to identify (or interpret) perceptual inputs. An implication of this approach is that theoretical progress can only be achieved by embedding theories of priming within theories of perception, and cognition more generally. Indeed, on this view, it is a mistake to develop a theory of long-term priming per se., but rather, theories of perception (semantics) need to be developed that support learning, with single-trial learning manifesting itself as priming in various tasks.

In the present chapter we highlight the role that functional considerations can play in improving our understanding of priming. We organise the chapter in four parts. First, we argue that priming is best explained as a by-product of learning within perceptual systems, with particular focus on visual word priming within the orthographic system. Although a number of authors have made similar functional claims, their theories have not been strongly constrained by these considerations, and as a consequence, have developed relatively independently of the large literature on word and object perception. Second, we consider an alternative approach, according to which priming and perception are embedded within a more general theory of memory – so-called instance theories of memory. Although this framework also emphasizes the adaptive function of the system that supports priming, we review evidence that we take to be problematic for this approach. Third, we outline an initial attempt to model visual word priming as an incidental by-product of learning within the orthographic system using a standard connectionist model of word recognition. Although limited in important ways, we argue that the model's success demonstrates the promise of this general approach. And fourth, we discuss how long-term priming techniques should be used as a tool to constrain theories of reading. The standard view in psycholinguists is that long-term priming is mediated by episodic rather than lexical processes, and as a consequence, it is considered an inappropriate tool to study visual word recognition processes. Indeed, masked priming techniques were introduced in order to eliminate these episodic effects (Forster & Davis, 1984). This characterization of long-term priming, in our view, is a mistake.

<u>Orthographic learning as the basis for visual word priming</u>.

The conclusion that long-term visual word priming is mediated by lexical-orthographic knowledge is based on numerous parallels between the structure of orthographic knowledge on one hand and priming results on the other. Consider the following three attributes of orthographic knowledge that have been proposed on the basis of various empirical results other than long-term priming. First, orthographic knowledge is coded in an abstract and modality specific format. So for instance, the visual patterns "r" and "R" map onto a common abstract letter identity r^*, which in turn map onto abstract lexical-orthographic codes, such that "read" and "READ" map onto an abstract word representation $read^*$ (e.g., Besner et al., 1984; Bowers, Vigliocco, & Haan, 1998; Coltheart, 1981; McClelland, 1976; McConkie & Zola, 1979). Second, orthographic knowledge encodes frequency information, with high frequency words identified more quickly than low frequency words (e.g., Forster & Chambers, 1973). And third, orthographic knowledge encodes morphological structure, such that the words CARS and CAR contact common orthographic codes, whereas the visually similar word CARD maps onto a separate lexical-orthographic representation (e.g., Caramazza, Laudanna, & Romani, 1988; Rapp, 1992).

Now, consider the following predictions that directly follow from the hypothesis visual word priming is mediated by orthographic representations – all of which have been supported. First, priming should be reduced for items spoken at study and written at test (or vice versa), or between pictures and the corresponding words, as in both cases, the orthographic representations are not fully accessed at study and test. This has been repeatedly obtained (e.g., Jacoby & Dallas, 1981; Rajaram & Roediger, 1993). Second, and for the same reason, little or no priming should be obtained between synonyms and translation equivalents in bilinguals since these items share different orthographic representations, and again, this is obtained (e.g., Durgunoglu, & Roediger, 1987; Roediger & Challis, 1992). Third, priming should be unaffected by study-to-test changes in the visual format of words when words are presented in familiar formats that allows the orthographic system to be engaged in the normal way. Although the story is somewhat complex, as discussed below, this is the standard result (Carr, Brown, & Charalambous, 1989; Clarke & Morton, 1983; Feustel, Shiffrin, & Salasoo, 1983; Scarborough, Cortese, & Scarborough, 1977; for review of these findings, see Bowers, 1996). Indeed, a pattern of abstract and modality specific priming is obtained even when the perceptual changes are substantial, for example, between visually dissimilar upper-and lower-case words in English <u>READ/read</u> (Bowers, 1996) and between Hiragana/Kanji scripts in Japanese (Bowers & Michita, 1998).

Forth, assuming that visual word priming is mediated by <u>lexical</u>-orthographic knowledge, priming should be constrained by morphological structure, such that priming is obtained between form related morphological relatives (e.g., cars/car), but not between items that are unrelated morphologically (e.g., card/car). Again, this is exactly what is obtained (e.g., Napps & Fowler, 1987). And fifth, assuming additive factor logic, priming should be sensitive to frequency, with more priming obtained for low- compared to high-frequency words. This is the case for the modality specific- and non-specific components of priming (Bowers, 2000b).

It is important to emphasize that all of these predictions follow directly from an orthographic account of long-term priming. This is not the case, however, for many memory based theories. For example, according to a transfer-appropriate-processing approach, priming is obtained to the extent that the same processes are engaged at study and test (e.g., Blaxton, 1989). Although this account can be reconciled with the above findings, it does so in an ad-hoc manner. In order to account for case-insensitive priming effects, for instance, it can be assumed that the visual processing of upper- and lower-case words is effectively the same, resulting in robust cross-case priming. But the theory could just as well predict priming to be sensitive to study-to-test changes in letter case, and indeed, early reports of case specific priming were taken as supportive of this approach (e.g., Roediger & Blaxton, 1987). Similarly, it is not at all clear that modality specific priming should be constrained by morphological status, although it can always be argued that similar visual processes are engaged when processing CARS and CAR, but not when processing CARD and CAR. The fact that the theory is consistent with a variety of possible outcomes highlights the lack of constraints on the theory.

As noted above, the hypothesis that priming is a by-product of learning or memory processes within orthographic and related systems is not new, and indeed, Morton (1979) and more recently Schacter (1990) and Ratcliff and McKoon (1997), and others have made similar claims. Nevertheless, relevant constraints that follow from this assumption are often ignored in the development of these theories. In the case of Morton (1979), priming was embedded in a model of word recognition that ignored issues of learning, and as a consequence, the model was hand-wired (fair enough, as it was developed over 20 years ago). On this model, identifying a word reduces the threshold of its corresponding logogen, and as a consequence, less information is required to identify the word when later presented—that is, priming is obtained. Reducing logogen thresholds is not a form of learning that increases the efficiency in perceiving repeated words, but rather, it lowers the standard of evidence required in order to identify repeated words. That is, priming is a product of bias rather than a change in sensitivity.

The issue of bias vs. sensitivity is currently an active topic of debate (see Bowers, 1999; McKoon & Ratcliff, in press; Ratcliff and McKoon, 1997), although there is growing consensus that priming does reflect, in part, a change in sensitivity (although, see Wagenmakers et al., chapter X). If indeed sensitivity does change as a consequence of priming, then Morton's account of priming would be falsified. But a more fundamental problem is that the model does not include any learning mechanisms that could support the acquisition of orthographic knowledge in the first place. Not only is this a limitation of the model as a model of word identification, but it also leads to difficulties in accommodating key priming phenomena; in particular, the finding that priming extends to various novel materials, including pseudowords (e.g., Bowers, 1994). Indeed, this is the main reason the model is so frequently rejected as a plausible theory of priming (e.g., Schacter, 1990).

Like Morton, Schacter (1990) has argued that visual word priming is the product of memory processes embedded within the orthographic system (what Schacter called a visual word form system). However, on this latter account, priming is assumed to improve the processing of repeated items (rather than cause bias), and priming is assumed to be mediated by new and highly specific memory traces encoded at study in addition to the possible contribution of abstract orthographic codes. The latter claim was based on the finding that priming extends to various unfamiliar materials, and that under some circumstances, is reduced following various changes in the format of words presented at study and test (e.g., Graf & Ryan, 1990). Schacter garners evidence from various neuropsychological populations in order to support this view (e.g., Schacter, Rapscak, Rubens, Tharan, & Laguna, 1990), and the related claim that that spoken word priming and visual object priming are mediated by auditory word form and structural description systems that support the identification of spoken words and objects, respectively (Schacter, 1992).

Schacter's approach is very much in line with our own. The main criticism we would raise is that Schacter has proposed a structural rather than processing theory of priming, and as a consequence, it is disconnected with processing theories of reading. In our view, the next key step is to explicitly link theories of priming with processing theories of word identification, such that priming is explained as a by-product of learning within these systems. In addition, we can't resist one quibble. That is, Schacter continues use of the terms "implicit memory" and "memory system" to describe long-term priming and the visual word form system, respectively. This in our view obscures the claim that priming mediated by a system whose primary function is visual word identification, not memory.

A more recent and computationally explicit account of visual word priming was proposed by Ratcliff and McKoon (1997); the counter model of priming. As in the Morton model, priming is understood within the context a processing model of word identification. And like the logogen model, it assumes that each word is represented by an abstract logogen like unit--so-called counters--and that word identification is achieved when the activation of a counter (measured in counts) passes some threshold. Furthermore, these counters are organized by similarity such that counters of similar words are close together within a cohort whereas counters of dissimilar words are far apart. As a consequence, similar items can be confused in the model, such that the presentation of <u>died</u> results in some counts being assigned to other items in its cohort, such as <u>lied</u>, especially under conditions in which <u>died</u> is flashed quickly so that the perceptual information is degraded.

Priming in the counter model is due to the fact that counters act as attractors. That is, exposure to a word at study causes the word's counter to attract a few counts more than it otherwise would, stealing them away from counters of other similar words. This attraction leads to a benefit in identifying repeated words since the counter “steals” counts that might otherwise have been mistakenly assigned to related words, and leads to costs when the study and test words differ, since the counter of the study word now steals counts away from the counter of the flashed word. In the original model (Ratcliff & McKoon, 1997), benefits equaled costs, and according, the model was a pure bias model of priming. But based on the more recent evidence that priming for low-frequency words is not all bias, McKoon and Ratcliff (in press) revised the model by changing a parameter so that priming for low-frequency words also reflects an improvement in processing. In both the original and the updated models, the authors also assumed that the attractive force that mediates priming only extends to words within cohort. So for example, the prior study of died would result in extra counts being stolen from lied, but not sofa. This can account for a counter-intuitive finding obtained in a forced choice task in which participants identify briefly flashed targets by selecting one of two alternative choice words. That is, bias is eliminated when the alternatives are dissimilar, such that studying died did not facilitate the identification of died when the alternatives are dissimilar, such as died and sofa. It is important to note, however, that relevant data on this latter point are mixed: Bowers (1999, unpublished) and Neaderhiser and Church (2000) have found priming under these conditions, whereas McKoon & Ratcliff (in press) and Masson (2000) do not. The basis of these discrepancies are unclear.

Regardless of the ultimate resolution of this latter dispute, there are problems with the updated theory. Although the model is computationally precise and can account for a wide variety of word priming data in perceptual identification tasks, the model was constructed around these data, and it is not clear to what extent it can explain other findings. And even within this domain, the failure to consider functional or empirical constraints from the psycholinguistic domain results in the model providing post-hoc rather than principled explanations of many of their findings. For example, as noted above, the authors initially obtained evidence that priming was pure bias, and accordingly, they did not include learning mechanisms in their model. When evidence was obtained that priming includes a change in sensitivity for low frequency words, the authors changed a parameter that supported learning for low frequency words. The latest version of the counter-model is now in the position in which it only includes learning mechanisms for pre-existing low-frequency words, and as a consequence, the model has to be hand-wired, much like the Morton (1979) model. It therefore inherits its chief weaknesses as a theory of priming, namely, it cannot support pseudoword priming.

In a similar way, the claim that priming is mediated by attractive forces between words within a cohort was solely based on their priming data. There is no obvious adaptive function for this claim, and the authors could change this parameter as well if subsequent work demonstrates that priming is obtained under conditions in which the alternatives are dissimilar. Indeed, because McKoon and Ratcliff (in press) now assume that priming for low frequency words is mediated by a change in sensitivity, the authors should expect priming for these items regardless of the similarity of the alternatives – contrary to their claim and their observations (although consistent with Bowers, 1999; Neaderhiser & Church, 2000). More generally, even within the limited scope of word identification in the perceptual identification task, without some principled constraints, the authors are free to vary parameters to accommodate whatever specific set of priming results are obtained.

According to Ratcliff and McKoon (1997), one of the key advantages of their theory over previous theories of priming is that it is implemented in a computational model. But contrary to the common claim, there is no intrinsic virtue of implementing a theory in computer-code, nor of modeling a restricted data set with great precision (Roberts & Pashler, 2000). Computational modeling is only a virtue when it allows predictions to be made that are too difficult to calculate without the aid of simulation. The crucial test of their theory is whether it can accommodate the large psycholinguistic database on visual word identification (making it a competitor to various other theories of visual word identification), and whether it makes novel and nontrivial predictions regarding priming and/or reading. Otherwise, the authors have made the mistake of building a model of visual word priming in a particular task rather than building a theory of word identification in which priming is a incidental by-product of learning (or bias).

In contrast with the above theories, we would argue that a number of predictions concerning visual word priming can be derived from a theory that is primarily designed around constraints associated reading written words. For example, there is now strong evidence that phonological codes are quickly (some would say automatically) activated following the visual presentation of words (e.g., Stone, Vanhoy, & VanOrden, 1997), and accordingly, we should predict phonological contributions to visual priming tasks. Indeed, phonological codes would provide an obvious candidate for cross-modal priming, and consistent with this hypothesis, priming effects obtained between homophones (which are presumably phonologically based) are approximately the same size as cross-modal priming (Rueckl, & Mathew, 1999; Ziemer & Bowers, submitted). Furthermore, a number of studies suggest that it is possible to manipulate the extent to which participants rely on phonological representations when responding in a lexical decision task: When nonword foils sound like words (e.g., brane), participants may be encouraged to rely on the orthographic rather than phonological representations (e.g., McQuade, 1981; Pugh, Rexer, & Katz, 1994; but see Pexman, Lupker, & Reggin submitted), whereas when participants are asked to categorize items in terms of their sound (such that brane is categorized as a word whereas blap would not), then participants must rely on phonological codes (e.g., Pexman et al., submitted). Accordingly, we might expect homophone and cross-modal priming to be reduced in the former condition and elevated in the latter case. Ziemer & Bowers (submitted) have verified the first prediction, and we (Kouider & Bowers) are currently testing the latter. In addition, various evidence suggests that orthographic and phonological representations are more strongly interconnected than visual object representations on the one hand and phonological and orthographic representations on the other (in fact, it is generally assumed that object and word knowledge are completely unconnected, with semantics intervening; e.g., Levelt, 1989). Accordingly, priming between pictures on the one hand and written or spoken words on the other should be reduced compared to cross-modal priming. Indeed, Rajaram & Roediger (1993) found that picture-to-written-word priming was reduced compared to spoken-to-written-word priming across four different tasks, and a similar finding is predicted when picture-to-spoken-word priming is compared to written-word-to-spoken-word priming. Numerous other predictions follow, and below, we describe a surprising prediction that has recently been confirmed. None of these predictions follow from standard memory based accounts, nor from previous theories that have embedded priming within processing models of word identification.

Instance theories of perception and priming.

There is another approach to understanding priming that also explicitly links theories of perception to theories of priming, but which is at odds with many of the claims made above. In contrast with the view that priming is best understood within the context of a more general theory of word and object perception, it is argued that theories of perception, semantics and priming should all embedded within a more general theory of episodic memory. On this approach, each experience is stored in detail as a separate memory trace, and the different ways this information is retrieved determines whether perception, semantics, episodic memory, or priming is expressed. On the strong version of this hypothesis, all of these diverse phenomena are seen as manifestation of different processes within a single episodic memory system. This obviously contrasts with the view described above in which written words are represented in terms of abstract logogen units, and where the orthographic system should be distinguished from related systems, such as phonological, semantic, and episodic systems. For a more detailed accounts of instance approaches, see the chapters by Whittlesea, this volume, as well as Goldinger 1998; Jacoby & Brooks, 1984; Hinztman, 1986, among others.

Despite the profound differences between “instance” accounts of priming and the “abstractionist” approach we’ve outlined above, they do share some fundamental properties. On both accounts visual word priming reflects a form of learning mediated by a system that supports visual word identification. On the instance view, however, learning (sometimes called skill acquisition within this framework) is mediated by the acquisition of new episodic instances within memory, making memory retrieval more efficient (Logan, 1990). In addition, because priming is mediated by a system that serves various functions, including object and word perception, both accounts are constrained by findings and theories outside the priming domain.

For example, if the abstractionist theory of object recognition advanced by Biederman, Hummel, and colleagues (Biederman, 1987; Hummel & Biederman, 1992) is correct, instance theories of priming would be falsified. On this account, objects are coded in memory in terms of abstract “structural descriptions” specifying an object’s parts in terms of categorical 3-D shape primitives (e.g., brick, cones, etc.) and their categorical relations to one another (e.g., on-top-of). So a table might be represented as a horizontal slab on top of four vertical posts. The abstract nature of these representations allows the model to categorize familiar and novel objects at a basic level since members of the category typically share the same description – i.e., most tables will be represented as a horizontal slab on top of four vertical posts, regardless of viewing angle. But as noted by critics of this approach, the categorical nature of these representations leads to problems when trying to distinguish between two different tables that share the same structural description (but see Hummel & Stankiewicz, 1998).

On the other hand, if "view based" models of object recognition are correct, abstractionist views of priming would be compromised. On this approach, objects are coded as holistic two-dimensional patterns as they appear from specific views, and identification consists in comparing holistic visual input patterns to these holistic memory representations (Tarr & Bulthoff, 1998). Because these visual codes represent the precise metrical information of the 2-D views, view based models show promise in explaining exemplar-specific object recognition (i.e., recognizing the same objects from different viewpoints), but they have not been tested extensively on basic level categorization. Indeed, according to Hummel (in press), this approach utterly fails at basic level categorization, as well as various other functions.

Similar conclusions follow from an ongoing debate regarding spoken word recognition. As detailed by Lachs et al. (chapter **), standard theories tend to assume that lexical-phonological codes are coded abstractly, and the task of word recognition is to discard the idiosyncratic acoustic details associated with a given spoken word in order to match it onto its stored abstract memory trace. The alternative is that words are stored in memory with all their details, and these details serve to facilitate identification rather than being regarded as noise. Again, a resolution to this debate speaks directly to theories of priming, and indeed, a wide range of non-priming phenomena provide key sources of constraint for theories of priming.

Despite the very different processing and representational assumptions of abstractionist and instance based approaches, it remains difficult to identify a critical experiment to distinguish between these views. Nevertheless, there are, in our view, empirical and theoretical reasons to prefer abstractionist accounts.

First, consider the priming data themselves. A number of authors claim that priming phenomena are sensitive to various study-to-test changes in the surface form of items, such as changing font or case of written words, or changing the voice of spoken words, and this is said to show that detailed perceptual memory traces support priming, and perception more generally. That is, the priming data are claimed to support instance based theories of cognition (e.g., Tenpenny, 1995). However, as argued in detail elsewhere (Bowers, 2000a), the data are more compatible with abstractionist approaches. The story is most straightforward in the visual domain, in which priming effects are often strikingly abstract. For example, using the perceptual identification task, Bowers (1996) assessed priming for words that are composed of letters that are visually dissimilar (A/a, B/b, D/d, E/e, G/g, L/l, and Q/q) in upper- and lowercase (e.g., DREAD/dread). Similar priming was obtained for study-test items presented in the same case (17% improvement over baseline) and different case (16% improvement), and at the same time, little priming was obtained when words were spoken at study (5% improvement). More striking, in two studies using the lexical decision task, similar priming was obtained for Japanese words studied and tested in the same script (Kanji-Kanji or Hiragana-Hiragana) or different script (Kanji-Hiragana or Hiragana-Kanji) despite the fact that the scripts are completely unrelated in the visual format (Bowers & Michita, 1998). Averaging across the two studies, same script priming was 28 ms and cross-script priming was 24 ms, a difference that did not approach significance. Once again, little priming was obtained when items were spoken at study (6 ms averaging across studies), suggesting that cross-script priming was mediated by abstract orthographic codes (also see Brown, Sharma, & Kirsner, 1984; Feldman & Moskovljevic, 1987). It is difficult to see how an instance based theory can account for modality-specific priming effects obtained between unrelated visual forms. And although advocates of instance based accounts have emphasized reports in which priming is reduced following study-to-test changes in visual format (e.g., Tenpenny, 1995), these effects are largely restricted to conditions in which the study or test words are presented in unusual formats (e.g., words presented upside down and mirror reversed). Given that these items are unlikely to be processed in the normal way within an orthographic system, abstractionist theories of word priming are not challenged by these findings (although the results do show that specific visual information is indeed coded in memory).

Priming results obtained with spoken words would seem to provide stronger evidence in support of instance based theories, as study-to-test perceptual changes more often influence priming (see Lach et al., this volume, for one characterization of this literature). But again, we would argue that the data are more compatible with abstractionist accounts. Luce and Lyon (1998), for instance, assessed spoken word priming in an auditory lexical decision task in which items were repeated in the same or different voice. The authors chose this task because test items are not degraded (as is required in perceptual identification, or completion tasks) and response latencies are brief (reducing any effects of explicit memory strategies). Under these conditions, no voice change effects were obtained despite robust priming effects. Recognition memory, by contrast, however, was sensitive to voice changes, showing again that perceptually specific information is encoded somewhere in memory. Consistent with this analysis, Schacter, Church and Bolton (1995) used an identification task in which words were degraded with a low-pass filter at test, and they observed voice specific priming in a group of participants with normal memory, whereas robust priming in the absence of voice specific effects was obtained in a group of amnesic patients. Accordingly, the voice effects in the normal group may have been due to explicit contamination. More recently, Pilotti et al. (2000) found spoken word priming to be completely insensitive to voice changes in the stem- and fragment-completion tasks, whereas priming was sensitive to voice change in an identification tasks that masked words with either noise of low-pass filtering. It is not clear why different results were obtained across tasks, but again, it shows that spoken word priming is often insensitive to large perceptual changes.

More recently, Pallier, Sebastian-Galles, and Colome (1999, submitted) have provided compelling evidence in support of an abstractionist approach using a different logic. The authors took advantage of the existence of two populations of Spanish-Catalan fluent bilinguals who share a common lexicon but have slightly different phonetic systems. In particular, Catalan-dominant speakers distinguish between certain phonemes (for instance, the consonantal contrast /s/- /z/) whereas Spanish-dominant bilinguals, even if they started learning Catalan as early as 6 years old, are unable do so. Participants performed lexical decisions to spoken Catalan words that were repeated verbatim or to words repeated with one phonetic feature changed (*casa-casa* versus *casa-caza*, meaning *'marry'* and *'house'*). The results revealed that the priming effects are modulated by abstract linguistic rather than the acoustic similarity. That is, the phonetic change eliminated priming for the Catalan-dominant speakers (they were perceived as different words), but for Spanish-dominant speakers, priming was equivalent to repetition priming (the two items were perceived as the same word).

To summarize, priming effects are often abstract, but specific visual and auditory information are encoded in memory and can affect priming and recognition memory under some conditions. What can be concluded from this? In deciding whether the data support abstractionist or instance theories, it is important to note that abstractionist theories do not reject the possibility that specific perceptual information is coded in memory – indeed, the information must be coded somewhere, as we can tell the difference between an upper- and lower-case letter, and perceive and remember differences between subordinate level items (Although in the case speech perception, phonetic changes can be difficult to detect. Indeed, Spanish-dominant bilinguals have a difficult time distinguishing between the phonemes /s/ and /z/ that are represented in Catalan but not in Spanish (Pallier, Bosch, & Sebastian-Galles, 1997). Thus, demonstrations of specific priming effects (or perceptually specific recognition memory) are not sufficient to rule out abstractionist views. It must also be shown that instances memories can support the various abstract results.

And indeed, advocates of instance theories claim that these theories can accommodate the abstract priming results (e.g., Tenpenny, 1995), as well as object recognition across changes orientation (Tarr & Bulthoff, 1998) and spoken word identification regardless of voice (e.g., Goldinger, 1998). But in the case of object recognition, there is almost no research concerned with how basic level identification can be achieved, and as noted above, Hummel (in press) argues that they fail on this basic function. And with regard to priming, it is not at all clear that these theories can support priming between visually dissimilar study/test items. In order to explain abstract priming within an instance framework, each input is thought to retrieve all similar memories stored in memory, creating an “echo” in which all instances blend together, producing a virtual abstract code in the absence of any abstract codes stored in permanent memory. This approach may indeed support abstract priming when the study-test words that differ in font, size, etc., because in these cases, there is some similarity between items so that the test item can retrieve the studied word amongst other related items. However, these approaches cannot accommodate modality specific priming effects between words that are arbitrarily related, such as READ/read (Bowers, 1996) and Hiragana and Kanji Japanese scripts (Bowers & Michita, 1998). The echo generated from a Kanji word at test will not be affected by the prior study of a Hiragana word, as there is no perceptual overlap. An echo could occur in semantics or phonology based on the similarity of these items in these domains, but it cannot account for the abstract priming results that are modality specific.

So how can abstractionist theories deal with the specific effects? A number of theorists have argued that the set of processes that support the effective categorization of items into basic level categories are incompatible with the goal of distinguishing different exemplars of the same category. On this view, different perceptual systems (or subsystems) support these two functions, so that both abstract and instance representations are encoded in memory separately (for different versions of this hypothesis, see Hummel & Stankiewicz, 1998; Marsolek, Kosslyn, & Squire, 1992.). Assuming both systems learn through experience, both should support priming. Specific priming effects would reflect priming within the specific system.

Of course, it might be objected that by advocating an abstractionist approach in which specific information can support priming makes the theory unfalsifiable. And indeed, this is a danger. The burden of this approach is to develop a theory in which the relation between abstract and specific information is well understood, allowing strong predictions to be made. Although this is not yet the case, it is interesting to note that there is some evidence that abstract and specific information is coded separately. For example, in the auditory modality, a number of studies indicate that patients with right hemisphere lesions show deficits in voice recognition (e.g., Van Lancker & Kreiman, 1987) and in processing prosodic aspects of speech (e.g., Coslett, Roeltgen, Rothi, & Heilman, 1987), whereas word identification disorders result from left hemisphere lesions (e.g., Schacter et al., 1993). Along the same lines, dichotic listening studies with non-brain damaged individuals have found a left-ear (right hemisphere) advantage for processing aspects of voice information, whereas identifying words (irrespective of voice) is best accomplished in the left hemisphere (e.g. Kimura, 1973). Further converging evidence has been reported using brain imaging technologies. Using PET, Zatorre, Evans, Meyer, and Gjedde (1992) reported that the left hemisphere is more active than the right when participants identified phonemes in spoken syllables (conditions in which surface details of the words were irrelevant), and the right hemisphere was more active than the left when participants made pitch discriminations to the same items (conditions in which surface details were critical). Schacter McGlynn, Milberg, & Church (1993) reported a patient with a large left-hemisphere lesion that led to <u>pure-word deafness</u>, such that he was able to repeat spoken words, but was unable to understand the meaning of these words. The patient showed robust priming for words repeated in the same voice (29% improvement over baseline) but minimal priming for words repeated in a different voice (4% improvement). One interpretation of this finding is that the patient's intact right hemisphere mediated his ability to repeat words and supported same-voice priming, but damage to abstract phonological codes in the left hemisphere (that normally serve as access codes to semantics) prevented him from understanding the meaning of spoken words, and reduced priming following the study/test voice change. Consistent with this hypothesis, control subjects showed much more abstract priming (23% same voice vs. 16% different voice), suggesting that the left hemisphere is indeed important for supporting priming following voice-changes (see Church & Schacter, 1994, for additional evidence in support of this view). So in all cases, abstract and specific information appear differentially lateralized. Marsolek and colleagues make similar claims for visual materials (e.g., Marsolek et al., 1992; but see Bowers & Turner, submitted). The story is complicated by the finding that abstract and specific information often interact in word processing tasks (e.g., see Lach et al., Chapter *, this volume), and these important findings will also need to be explained. Also, there is some evidence that attention can modulate the degree of abstraction, as Stankiewicz, Hummel, and Cooper (1998) only found abstract priming between mirror inverted pictures when study items were attended, whereas repetition priming did not require attention. Despite these complications, as long as abstract information is encoded somewhere in perceptual (and semantic) systems, instance theories are falsified.

A second argument in support of abstractionist views is more fundamental. One of the virtues positing abstract (and localized) mental representations is that they can be incorporated into generative systems that allow the construction of complex mental representations from simple representations—that is, they can support <u>compositional</u> representations (Fodor & Pylyshyn, 1988). This in turn allows visual, phonological, and lexical-semantic (that is, linguistic) systems to represent information that extends beyond past experience. This is not to say that abstract representations are sufficient to support these functions, but they play a necessary role in neural networks models of object identification (e.g., Hummel & Biederman, 1992) and semantics (e.g. Hummel & Holyoak 1997) that implement these functions. By contrast, it is not clear how generative systems can be supported in an instance model of memory. Thus far, instance models have largely been applied to remembering or perceiving single words, with no consideration of how virtual abstract codes (e.g., echoes) could be fed into a language processor that must construct complex representations from simpler parts. Of course, it is not reasonable to rule out instance models because they currently fail to realize these high-level objectives, but it should at least be possible to <u>imagine</u> how such a model could scale up to achieve these functions. But thus far, there are not even speculations.

Finally, as a more general point, it is often claimed that instance models are more parsimonious, as they do away with multiple memory systems in favor of a single episodic system that serves various functions, from visual and spoken word identification, episodic memory, etc.. And indeed, we agree that it is not helpful to attribute these different functions to separate <u>memory</u> systems. But in our view, the claim that all these functions are mediated by a single system loses sight of some basic distinctions. Not only do vision, audition, semantics, episodic memory deal with different types of information, they serve different functions, and they need to process information in qualitatively different ways. For example, in the case of vision, information is distributed over space, whereas in audition, information is distributed across time. It seems unlikely, at least to us, that these different demands are performed by a single system. To describe these as different processes within a single system amounts to the claim that we have different processes within a single brain – certainly true, but not particularly informative.

Modeling visual word priming within a theory of word identification that supports learning.

In order to lend some credibility to the claim that word priming is an incidental by-product of learning processes embedded within system(s) whose function is to identify words, Bowers, Damian, and Havelka (2000) attempted to simulate priming in the distributed, developmental model of word recognition and naming put forward by Seidenberg and McClelland (1989). At first, this might seem an odd choice given the above discussion: The model does not include abstract orthographic representations that could cross-case or morphological priming (although connectionist models have been developed that can learn abstract letter codes, Polk & Farah, 1997; and morphological relations, Plaut & Gonnerman, in press), and the model does not include localist representations that could be fed into a generative system (although connectionist models with localist representations can be developed with this properties as well, Hummel & Holyoak, 1997). Indeed, as noted by Bowers (2000c) models with distributed representations cannot even represent two things at a time, despite the obvious need to do so in semantic, phonological, and related systems.

However, we were not interested in evaluating the merits of this model as a model of reading, or as a model of priming. Rather, we were simply interested as to whether a model that learns by back-propagation—a learning rule common to many connectionist models--could in principle support long-term priming following a single study episode. Surprisingly, word priming has not previously been simulated within a model trained for reading words, and it was not at all obvious to us that they would be successful. Learning with back-propagation is subject to a phenomenon called the stability-plasticity problem (often called “catastrophic interference”) in which new information erases old information (Grossberg, 1987; McClosky & Cohen, 1989). In order to avoid this interference so that large vocabularies can be acquired, it has been necessary to reduce the learning rates, and to introduce an interleaved study regimen in which all the words in the vocabulary are acquired in parallel. This solution works fine for some purposes, but it raises the question as to whether this slow learning can support long-term priming that occurs following a single-study trial. Indeed, Ratcliff and McKoon (1997) argued that the Seidenberg and McClelland (1989) model was incapable of supporting long-term priming because of its slow learning rates. And even if the learning rates are found to be sufficient to support priming for a short duration, the interference effects due to learning unrelated words processed between study and test may be incompatible with the longevity of priming.

We used the fully trained Seidenberg and McClelland (1989) model in order to simulate word priming, without changing any of its parameters. In order to follow standard procedures in behavioral studies, we ran the simulations numerous times for each experiment, with each simulation corresponding to a single participant. In the study phase of each simulation, the model was presented with the study words, and the connection weights were modified to the same degree as the learning trials on which it was originally trained. That is, the study words were simply the most recent words the model was trained on. As in behavioral studies, study words were rotated through the various conditions, and at test, all target words were presented. Priming was computed by comparing the orthographic error scores for words in the repeated and baseline conditions, averaging across all simulations. These error scores were then translated to estimates of the reaction time using a formula described by Seidenberg and McClelland (1989, p. 532). Note, the learning these estimates were not derived to fit long-term priming data, but were based on various naming and lexical decision studies the authors reported. Nevertheless, robust priming effects were obtained in the model, and the pattern of priming mirrored priming in various experiments.

For example, priming was greatly reduced for a set of high- compared to low-frequency words, typical in experimental studies (e.g., Bowers, 2000b). Furthermore, these priming effect persisted when 1000s of unrelated learning trials were interposed between study and test, indicating that the priming effects were long-lasting. In addition, and contrary to our expectations, priming effects were largely word specific. In our simulation study, a robust repetition- contrasted with small form-related (e.g., prime = <u>boast</u>, target = <u>toast</u>) priming effect, consistent with a large number of behavioral studies (e.g., Napps & Fowler, 1987; Ratcliff & McKoon, 1997). In trying to understand why a model with distributed representations only produced small form-priming effects, we discovered that form priming in the model was slightly larger for pairs that rhymed, and slightly inhibitory for pairs that did not rhyme. We subsequently ran a behavioral study (Bowers et al., 2000), and found a similar pattern. So here, the model made a novel and counter-intuitive prediction that was subsequently confirmed.

The finding that facilitation and cost is obtained in the model for the rhyme and non-rhyme form related words has another implication. In a number of recent reports, it has been claimed that changes in sensitivity result in no costs, for example, Keane, Verfaellie, Gabrieli, and Wong (2000) wrote “According to a sensitivity account, by contrast, only the benefit, and not the cost, should be observed: If priming improves the ability to extract perceptual information from a stimulus, then identification of a word should be enhanced by prior exposure to that word and should not be harmed by prior exposure to its orthographic mate” (p. 318), and it has been proposed that separate mechanisms underlie benefits and costs in priming (e.g., McKoon & Ratcliff, in press). Although the latter claim may turn out to be true, the present findings highlight the fact that both benefits and costs can be the product of a single learning system whose function is to improve processing. A particularly compelling example of costs associated with learning can be found in a related language domain. At a few months, babies are able to perceive and produce the phonemes of all the languages of the world, but as they are exposed and learn the phonology of their particular language, they become insensitive to key phonetic distinctions in other languages while becoming more adept in identifying the phonemes within their own language (Kuhl et al., 1992). A familiar example is that Japanese speakers have great difficulty in perceiving and producing the phonemes “l” and “r” in English (Goto, 1971). It seems unlikely that a separate bias mechanisms underlies this deficit, just as there are no bias mechanisms underlying the costs we observed in these simulations. Learning systems that improve processing do not rule out cost in performance, although the benefits outweigh costs in the domain in which they function.

So in sum, a connectionist model that learns to pronounce individual words not only supports various long-term priming effects, but it also made a novel prediction that was confirmed experimentally. Although we do not take these findings as providing support for this particular model of word identification, or indeed, models that learn distributed representations using back-propagation, we do take these findings as providing general support for the view that visual word priming is best understood within the context of a model that is designed to identify words. On this view, theoretical advances in theories of word identification will improve our understanding of priming, and priming may provide important data concerning the types of learning that should be incorporated in models of reading.

Using long-term priming as a tool to address psycholinguistic questions.

Although we think that there are compelling reasons to view long-term word priming as a by-product of learning within the orthographic system, this phenomenon is rarely used by psycholinguists interested in word identification processes. Instead, researchers interested in visual word recognition have relied heavily on masked priming paradigms in which a pattern mask (e.g., ######) is replaced by a prime word that is briefly flashed (e.g., 50 ms), which in turn is replaced by the target. Under these conditions, primes are typically unnoticed by the participants, but these items nevertheless facilitate processing of targets when prime and target are the same compared to different.

Although masked and long-term priming share many features (e.g., Bowers, 2000a), three particular attributes of long-term priming are often cited as reasons to reject this procedure as a useful psycholinguistic technique. First, Forster and Davis (1984) noted that long-term priming is reduced for high frequency compared to low frequency words, what the authors called the frequency attenuation effect. By contrast, a frequency attenuation effect generally does not occur in masked priming (but see Bodner & Masson, 1997). Given that recognition memory is generally better for low-frequency words (e.g., Balota & Neely, 1980), the discovery of this qualitative difference led Forster & Davis (1984) to claim that “…long-term priming effect is totally mediated by episodic factors, whereas the short-term effect is an automatic consequence of repeated access of the same lexical entry." (p.694). The masking procedure was presumed to eliminate this episodic involvement in priming, suppressing the frequency attenuation effect. If in fact the authors are correct, then it follows that long-term priming is not a suitable technique to study abstract orthographic codes (unless of course episodic memories support lexical access, as assumed by instance theorists).

However, there are problems with this conclusion. Forster and Davis have considered the frequency attenuation effect to be an indicator of episodic influences whereas several studies have provided a dissociation between episodic recognition and the frequency attenuation effect (Rajaram & Neely, 1992; Duchek & Neely, 1989; Scarborough et al., 1977). For instance, Rajaram and Neely (1992) addressed this issue by using a modified episodic recognition task in which stimuli are presented under conditions identical to a primed lexical decision task. Whereas they found the classical frequency attenuation effect for long-term priming, they found an advantage for high-frequency words in the recognition task, exactly the opposite to priming. Rajaram and Neely concluded that these data pose a serious challenge to the Forster and Davis claim. Further evidences for a dissociation between long-term priming and recognition memory effects comes from the memory literature. For instance, densely amnesic patients show robust (sometimes normal) priming despite poor (sometimes chance) performance on explicit recall and recognition tests (Warrington & Weiskrantz, 1974). These dissociations between priming and episodic recognition memory suggests that the Forster & Davis inference about long-term priming was premature. And as described above, the simulation study by Bowers et al. (2000) demonstrated that a frequency attenuation effect can be produced by learning processes embedded within an orthographic system.

A second criticism of long-term put forward by Forster and Davis (1984) and others (e.g., De Groot & Nas, 1991) is based on findings reported by Oliphant (1983). Oliphant contrasted repetition effects in two conditions. In the first, participants performed a lexical decision both to the prime and target (as is common in this paradigm). In the second, they read prime words embedded in running sentences during the instruction session and then at test performed lexical decisions to repeated target words. Oliphant failed to obtain priming in this latter condition and concluded that participants have to be aware that target words are repeated for priming to occur. This awareness was assumed to result from overt responses on prime words and was supposed to be eliminated (along with priming) when words were embedded in text, leading him to support an episodic account.

Note, however, that a reduction in priming for words presented in the context of sentences is not surprising from an abstractionist position that attributes priming to the strengthening of abstract orthographic codes. Under these condition, there is no guarantee that prime words are perceptually processed to the same extent as when they are presented in isolation. Indeed, there is no reason to presume that structural changes occur automatically each time a word is identified, and there are many circumstances in which it is plausible to assume that the relevant changes do not take place. For example, when reading text, people tend to fixate on each specific word for only a brief moment, with an average fixation of approximately 200-250 ms (Rayner & Sereno, 1994). This may not be sufficient time to support the long-term modification of orthographic codes, particularly when multiple words are processed in running sentences. Consistent with this general analysis, Subramaniam, Biederman and Madigan (2000) recently reported a failure to obtain any priming for pictures presented up to 31 times in an RSVP sequence when items were displayed for between 72-126 ms/picture. At the same time, the pictures could be identified at these durations, and accordingly, the authors argued that priming requires participants to attend to an item for a period of time after the item has been identified. But in any case, Oliphant (1983) assessed priming for high-frequency words. When MacLeod (1989) replicated Oliphant's (1983) study using lower frequency words, priming was obtained.

A third criticism is that long-term priming often extends to nonwords (e.g., Besner & Swan,1982; Kirsner & Smith, 1974 ; Scarborough et al., 1977) whereas nonword priming is reduced or eliminated in the masked priming procedure introduced by Forster and Davis (1984) in which participants make lexical decisions to targets. The finding that long-term priming extends to nonwords is often taken as evidence that episodic memory mediates these effects, and masking the prime is presumed to eliminate episodic involvement, reducing the nonword priming. However, subsequent work has found a close parallel between the masked and long-term priming for nonwords. In particular, when participants make lexical decisions to the targets, nonword priming tends to be absent in both long-term and masked priming procedures (for long-term results, see Bentin & Moscovitch, 1988; Bowers, 1994; Fowler, Napps & Feldman, 1985; for masked results, see Forster, 1998; Forster & Davis, 1984), although exceptions are obtained in both tasks (for long-term results, see Besner & Swan,1982 ; Kirsner & Smith, 1974 ; Scarborough et al., 1977; for masked results, see Bodner & Masson, 1997; Sereno, 1991). As argued by Feustel, Shiffrin & Salasoo (1983) and by Humphreys, Evett, and Quinlan (1990), the failure to obtain nonword priming in these two paradigms is due to an idiosyncratic property of the lexical decision task. That is, subjects have a bias to respond "word" to the repeated nonwords since familiar letter strings tend to be words. This bias counteracts the improved perceptual processing of the nonwords, resulting in no priming. Consistent with this argument, when this response bias is eliminated by using an identification task, nonword priming effects are consistently obtained in both long-term and masked priming tasks (for long-term priming, see Bowers, 1994; Carr, Brown, & Charalambous, 1989; Feustel, Shiffrin & Salasoo, 1983; Kirsner & Smith, 1974; Rueckl, 1990; Salasoo, Shiffrin & Feustel, 1985; Whittlesea & Cantwell, 1987; for masked priming, see Humphreys et al., 1990; Masson & Isaak, 1999, among others). Accordingly, there is no reason to prefer masked over long-term priming techniques based on the nonword results.

Despite these considerations, the original arguments of Forster and Davis (1984) continue to be widely accepted in psycholinguistic circles (e.g., Frost, Deutsch, Gilboa, Tannenbaum & Marslen-Wilson, in press). The irony of this is that long-term priming has been one the major source for the development of single-word processing theories. In particular, Morton's logogen model has been one of the most influential models of word recognition, and its development was mainly driven by long-term priming data. Consider the evolution of the model from its first 1969 version until its final and more elaborated 1980 version. Based on long-term priming data showing reduced cross-modal repetition between words and pictures or between written and spoken words, the model was constrained to include different modality-specific logogen systems. Although not all current models include logogens, the majority of current activation models of visual word recognition include these modality-specific distinctions (see, for instance, Coltheart, Rastle, Perry, Langdon & Ziegler, in press, for an historical review of how the logogen model evolved and how its last version influenced the DRC model). On our view, long-term priming techniques can contribute once again to the development these models.

Note, we do agree with Forster and Davis that masked and long-term priming are supported by two different mechanism but, contrary to their conclusion, we claim these different mechanisms act onto the *same* abstract orthographic representations. Regarding masked priming, pre-existing abstract lexical codes are temporally activated (or opened) with no systemic structural change resulting. Thus, this form of priming has only short-term consequences and for those reasons it lasts only a few seconds. But on our view, long-term priming reflects structural changes (learning) in the orthographic system that affect the later processing of the repeated items, with larger improvements associated with lower-frequency words (e.g., Bowers, 2000b). On this later view, long-term priming should provide constraints to theories of word recognition, particularly those concerned with issue of learning.

Summary

Although long-term priming has been extensively studied during the past 20 years, there are relatively few detailed theories of priming, and even fewer theories that have paid careful attention to the underlying function of priming. We would suggest the word priming is best understood as a by-product of learning within orthographic and phonological systems, and as such, it is a mistake to develop a theory of priming per se. Rather, it is important to develop theories of perception (and perhaps semantics) that learn. Such an approach leads to novel and counter-intuitive predictions, some of which we have already confirmed.

References

Balota, D.A., & Neely, J.H. (1980). Test-expectancy and word-frequency effects in recall and recognition. Journal of Experimental Psychology: Human Learning and Memory, 6, 576-587.

Bentin, S., & Moscovitch, M. (1988). The Time Course Of Repetition Effects For Words and Unfamiliar Faces. Journal of Experimental Psychology-General, 117, 148-160.

Besner, D., Coltheart, M. & Davelaar, E. (1984). Basic Processes in reading: Computation of abstract letter identities. Canadian Journal of Psychology, 38, 126-134.

Besner, D., & Swan, M. (1982). Models Of Lexical Access In Visual Word Recognition. Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology, 34, 313-325.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. Psychological Review, 94, 115-147.

Blaxton, T.A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. Journal of Experimental Psychology: Leaning, Memory, & Cognition, 15, 657-668.

Bodner, G.E., & Masson, M.E.J. (1997). Masked repetition priming of words and nonwords: Evidence for a nonlexical basis for priming. Journal of Memory and Language, 37, 268-293.

Bowers, J.S. (1994). Does implicit memory extend to legal and illegal nonwords? Journal of Experimental Psychology: Learning, Memory, and Cognition, 20. 534-549.

Bowers, J.S. (1996). Different perceptual codes support word and pseudoword priming: Was Morton right all along? Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 1336-1353.

Bowers, J. S. (1999). Priming is not all bias: Commentary on Ratcliff and McKoon (1997). Psychological Review, 106, 582-596.

Bowers, J. S. (2000a). In defense of abstractionist theories of word identification and repetition priming. Psychonomic Bulletin & Review, 7, 83-99

Bowers, J. S. (2000b). The modality specific and non-specific components of long-term priming are frequency sensitive. Memory & Cognition(28), 406-414.

Bowers, J.S. (2000c). Challenging the widespread assumption that connectionism and distributed representations go hand-in-hand. Manuscript submitted for publication.

Bowers, J.S., Arguin, M. & Bub, D. (1996). Fast and specific access to orthographic knowledge in a case of letter-by-letter surface alexia. Cognitive Neuropsychology, 13, 525-567.

Bowers, J.S. & Damian, M, & Havelka, J. (2000). Can distributed orthographic knowledge support word specific long-term priming? Apparently so. Manuscript submitted for publication.

Bowers, J .S., & Michita, Y. (1998). An investigation into the structure and acquisition of orthographic knowledge: Evidence from cross-script Kanji-Hiragana priming. Psychonomic Bulletin & Review, 5, 259-264.

Bowers, J.S., & Turner, E.L. (2000). In search of abstract and specific visual form systems in the left and right hemispheres. Manuscsript submitted for publication

Bowers, J.S., Vigliocco, G., & Haan, R. (1998). Orthographic, phonological, and articulatory contributions to masked letter and word priming. Journal of Experimental Psychology: Human Perception and Performance, 24, 1705-1719.

Brown, H.L., Sharma, N.K., & Kirsner, K. (1984). The role of script and phonology in lexical representation. Quarterly Journal of Experimental Psychology: 36(A), 491-505.

Caramazza, A., Laudanna, A., & Romani, C. (1988). Lexical Access and Inflectional Morphology. Cognition, 28, 297-332.

Carr, T. H., Brown, J. S., & Charalambous, A. (1989). Repetition and Reading - Perceptual Encoding Mechanisms Are Very Abstract But Not Very Interactive. Journal of Experimental Psychology-Learning Memory and Cognition, 15, 763-778.

Church, B. A., & Schacter, D. L. (1994). Perceptual Specificity Of Auditory Priming - Implicit Memory For Voice Intonation and Fundamental-Frequency. Journal of Experimental Psychology-Learning Memory and Cognition, 20, 521-533.

Clarke, R. G. B., & Morton, J. (1983). Cross modality facilitation in tachistoscopic word recognition. Quarterly Journal of Experimental Psychology, 35A, 79-96.

Coltheart, M. (1981). Disorders of reading and their implications for models of normal reading. Visible Language, 3, 245-286.

Coltheart, M., Rastle, K., Perry, C., Langdon, R. & Ziegler, J. (in press). DRC: A Dual Route Cascaded model of visual word recognition and reading aloud. Psychological Review.

Coslett, H.B., Roeltgen, D.P., Rothi, L.G., & Heilman, K.M. (1987). Transcortical sensory aphasia: Evidence for subtypes. Brain and Language, 32, 362-378., 1987.

Davis, C. (1999). The Self-Organising Lexical Acquisition and Recognition (SOLAR) model of visual word recognition. Unpublished doctoral dissertation, University of New South Wales.

De Groot, A.M.B., & Nas, G.L.J. (1991). Lexical representation of cognates and noncognates in compound bilinguals. Journal of Memory and Language, 30, 90-123.

Duchek JM, Neely JH. (1989). A dissociative word-frequency X levels-of-processing interaction in episodic recognition and lexical decision tasks. Memory and Cognition, 17, 148-62.

Durgunoglu, A. Y., & Roediger, H. L. (1987). Test Differences In Accessing Bilingual Memory. Journal of Memory and Language, 26(4), 377-391.

Feldman, L. B., & Moskovljevic, J. (1987). Repetition priming is not purely episodic in origin. Journal of Experimental Psychology: Learning, Memory and Cognition, 13, 573-581.

Feustel, T. C., Shiffrin, R. M., & Salasoo, A. (1983). Episodic and Lexical Contributions to the Repetition Effect In Word Identification. Journal of Experimental Psychology-General, 112(3), 309-346.

Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. Cognition, 28, 3-71.

Forster, K. I. (1998). The pros and cons of masked priming. Journal of Psycholinguistic Research, 27(2), 203-233.

Forster, K. I., & Chambers, S. M. (1973). Lexical access and naming time. Journal of Verbal Learning and Verbal Behavior, 12, 627-635.

Forster, K.I. & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,680-698.

Fowler, C. A., Napps, S. E., & Feldman, L. (1985). Relations Among Regular and Irregular Morphologically Related Words In the Lexicon As Revealed By Repetition Priming. Memory & Cognition, 13(3), 241-255.

Frost, R., Deutsch, A., Gilboa O., Tannenbaum, M., & Marslen-Wilson, W. (in press). Morphological priming: Dissociation of phonological, semantic, and morphological factors. Memory and Cognition.

Goldinger, S.D. (1998). Echoes of echoes? An episodic theory of lexical access. Psychological Review, 105, 251-279.

Goto, H., (1971). Auditory perception by normal Japanese adults of the sounds `l' and `r'. Neuropsychologia, 9, 317-323.

Graf, P., & Ryan, L. (1990). Transfer-Appropriate Processing For Implicit and Explicit Memory. Journal of Experimental Psychology-Learning Memory and Cognition, 16(6), 978-992.

Grossberg, S. (1987). Competitive Learning - From Interactive Activation to Adaptive Resonance. Cognitive Science, 11, 23-63.

Hintzman, D.L. (1986). Schema abstraction in a multiple trace model. Psychological Review, 93, 411-428.

Hummel, J. E., & Biederman, I. (1992). Dynamic Binding In a Neural Network For Shape-Recognition. Psychological Review, 99(3), 480-517.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. Psychological Review, 104(3), 427-466.

Hummel, J. E., & Stankiewicz, B. J. (1998). Two roles for attention in shape perception: A structural description model of visual scrutiny. Visual Cognition, 5(1-2), 49-79.

Hummel, J.E. (in press). Where view-based theories break down: The role of structure in shape perception and object recognition. To appear in E. Dietrich and A. Markman (Eds.). Cognitive Dynamics: Conceptual Change in Humans and Machines., Cambridge, MA: MIT Press.

Humphreys, G. W., Evett, L. J., & Quinlan, P. T. (1990). Orthographic Processing In Visual Word Identification. Cognitive Psychology, 22(4), 517-560.

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic Cognition - Memory, Perception, and Concept-Learning. Psychology Of Learning and Motivation-Advances In Research and Theory, 18, 1-47.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. Journal of Experimental Psychology: General, 110, 306-340.

Keane, M. M., Verfaellie, M., Gabrieli, J. D. E., & Wong, B. M. (2000). Bias effects in perceptual identification: A neuropsychological investigation of the role of explicit memory. Journal of Memory and Language, 43, 316-344.

Kimura, D. (1973). The asymmetry of the human brain. Scientific American, 228, 70-78.

Kirsner, K., & Smith, M. C. (1974). Modality effects in word identification. Memory & Cognition, 2, 637-640.

Kirsner, K., Smith, M.C., Lockhart, R.S., King, M.L., & Jain, M. (1984). The bilingual lexicon: Language-specific units in an integrated network. Journal of Verbal Learning and Verbal Behavior, 23, 519-539.

Kuhl, P. K., Williams, K. A., Lacerda, F., Stevens, K. N., & Lindblom, B. (1992). Linguistic experiences alter phonetic perception in infants by 6 months of age. Science, 255, 606-608.

Levelt, W. J. M. (1989). Speaking: From Intetion to Articulation. Cambridge, MA: MIT Press.

Logan, G.D. (1990). Repetition priming and automaticity: Common underlying mechanisms? Cognitive Psychology, 22, 1-35.

Luce, P. A., & Lyons, E. A. (1998). Specificity of memory representations for spoken words. Memory & Cognition, 26(4), 708-715.

MacLeod, C.M. (1989). Word context during intial exposure influences degree of priming in word fragment completion. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 398-406.

Marsolek, C.J., Kosslyn, S.M. & Squire, L.R. (1992). Form-specific visual priming in the right cerebral hemisphere. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 492-508.

Masson, M. E. J (2000). Bias Effects in Repetition Priming of Masked Word Identification: Differential Influences of Modality. Manuscript submitted for publication.

Masson, M. E. J., & Isaak, M. I. (1999). Masked priming of words and nonwords in a naming task: Further evidence for a nonlexical basis for priming. Memory & Cognition, 27, 399-412.

McClelland, J.L. (1976). Preliminary letter identification in the perception of words and nonwords. Journal of Experimental psychology: Human Perception and Performance, 3, 80-91.

McClosky, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bowers (Ed.), The psychology of learning and motivation . New York: Academic Press.

McConkie, G. W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading. Perception & Psychophysics, 25, 221-224.

McKoon, G., & Ratcliff, R. (in press). The counter model for word identiciation: Reply to Bowers (1999). Psychological Review.

McQuade, D. V. (1981). Variable reliance on phonological information in visual word recognition. Language and Speech, 24, 99-109.

Morton, J. (1979). Facilitation in word recognition: Experiments causing change in the logogen model. In P. A. Kolers, M. E. Wrolstad, & H. Bouma, Processing models of visible language. New York: Plenum.

Napps, S.E., & Fowler, C.A. (1987). Formal relationships among words and the organization of the mental lexicon. Journal of Psycholinguistic Research, 16, 257-272.

Neaderhiser, B. J. & Church, B. A. (2000, April). Priming in forced choice identification: Memory or bias? Paper presented at the 92nd Meeting of the Southern Society for Philosophy and Psychology, Atlanta.

Oliphant, G. W. (1983). Repetition and recency effects in word recognition. Australian Journal of Psychology, 35, 393-403.

Pallier, C., Bosch, L., & Sebastian-Galles, N. (1997). A limit on behavioral plasticity in speech perception. Cognition, 64, B9-B17.

Pallier, C., Sebastian-Gallès, N., & Colomé, A. (1999). Phonological representations and repetition priming. Proceedings of Eurospeech '99, Budapest, Hungary, Sept. 5-9, 1999., vol. 4 (pp. 1907-1910).

Pallier, C., Sebastian-Gallès, N., & Colomé, A. (submitted). Phonological representations and repetition priming.

Pexman, P. M., Lupker, S. J., & Reggin, L. D. (submitted). Investigating the impact of feedback phonology in visual word recognition. .

Pilotti, M., Bergman, E. T., Gallo, D. A., Sommers, M., & Roediger III, H. L. (2000). Direct comparison of auditory implcit memory tets. Psychonomic Bulletin & Review, 7, 347-353.

Plaut, D. C. and Gonnerman, L. M. (in press) Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? Language and Cognitive Processes.

Polk, T.A., & Farah, M.J. (1997). A simple common contexts explanation for the development of abstract letter identities. Neural Computation, 9, 1277-1289.

Pugh, K. R., Rexer, K., & Katz, L. (1994). Evidence of flexible coding in visual word recognition. Journal of Experimental Psychology: Human Perception and Performance, 20, 807-825.

Rajaram, S., & Neely, J.H. (1992). Dissociative masked repetition priming and word-frequency effects in lexical decision and episodic recognition tasks. Journal of Memory and Language, 31, 152-182.

Rajaram, S. & Roediger, H.H. III (1993). Direct comparison of four implicit memory tests. Journal of Experimental Psychology: Learning, Memory, and Cognition, 19, 765-776.

Rapp, B. C. (1992). The Nature Of Sublexical Orthographic Organization - the Bigram Trough Hypothesis Examined. Journal of Memory and Language, 31, 33-53.

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

Rayner, K., & Sereno, S.C., (1994). Eye movements in reading: Psycholinguistic studies. In M.A., Gernsbacher (Ed.). Handbook of psycholinguistics. New York: Academic Press.

Roberts, S. & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing in psychology. Psychological Review.

Roediger, H. L., & Blaxton, T. A. (1987). Effects Of Varying Modality, Surface-Features, and Retention Interval On Priming In Word-Fragment Completion. Memory & Cognition, 15(5), 379-388.

Roediger, H. L., & Challis, B. H. (1992). Effects Of Exact Repetition and Conceptual Repetition On Free-Recall and Primed Word-Fragment Completion. Journal of Experimental Psychology-Learning Memory and Cognition, 18(1), 3-14.

Roediger, H.L. III, & McDermott, K.B. (1993). Implicit memory in normal human subjects. In F. Boller & J. Grafman (Eds.), Handbook of neuropsychology (Vol. 8, pp. 63-131). New York: Elsevier.

Rueckl, J. G. (1990). Similarity Effects In Word and Pseudoword Repetition Priming. Journal of Experimental Psychology-Learning Memory and Cognition, 16, 374-391.

Rueckl, J. G., & Mathew, S. (1999). Implicit memory for phonological processes in visual stem completion. Memory & Cognition, 27, 1-11.

Scarborough, D. L., Cortese, C., & Scarborough, H. S. (1977). Frequency and repetition effects in lexical memory. Journal of Experimental Psychology: Human Perception and Performance, 3, 1-17.

Schacter, D. L. (1990). Perceptual representation system and implicit memory: Toward a resolution of the multiple memory debate. In A. Diamond (Ed.), Development and Neural Basis of Higher Cognitive Function. Annals of the New York Academy of Science.

Schacter, D. L. (1992). Priming and Multiple Memory-Systems - Perceptual Mechanisms Of Implicit Memory. Journal of Cognitive Neuroscience, 4(3), 244-256.

Schacter, D. L., Church, B., & Bolton, E. (1995). Implicit Memory In Amnesic Patients - Impairment Of Voice-Specific Priming. Psychological Science, 6(1), 20-25.

Schacter, D.L., McGlynn, S.M., Milberg, W.P., & Church, B.A. (1993). Spared priming despite impaired comprehension: Implicit memory in a case of word meaning deafness. Neuropsychology, 7, 107- 118.

Schacter, D. L., Rapscak, S. Z., Rubens, A. B., Tharan, M., & Laguna, J. (1990). Priming Effects In a Letter-By-Letter Reader Depend Upon Access to the Word Form System. Neuropsychologia, 28, 1079-1094.

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

Southern Society for Philosophy and Psychology, Atlanta.

Sereno, J. A. (1991). Graphemic, associative, and syntactic priming effects at brief stimulus onset asynchrony in lexical decision and naming. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 459-477.

Stankiewicz, B. J., Hummel, J. E., & Cooper, E. E. (1998). The role of attention in priming for left-right reflections of object images: Evidence for a dual representation of object shape. Journal of Experimental Psychology-Human Perception and Performance, 24, 732-744.

Stone, G. O., Vanhoy, M., & VanOrden, G. C. (1997). Perception is a two-way street: Feedforward and feedback phonology in visual word recognition. Journal of Memory and Language, 36(3), 337-359.

Subramaniam, S., Biederman, I., & Madigan, S. A. (2000). Accurate identification but no priming and chance recognition memory for pictures in RSVP sequences. Visual Cognition, 7, 511-535.

Tarr, M. J., & Bulthoff, H. H. (1998). Image-based object recognition in man, monkey and machine. Cognition, 67(1-2), 1-20.

Tenpenny, P.L. (1995). Abstractionist versus episodic theories of repetition priming and word identification. Psychonomic Bulletin & Review, 2 339-363.

Van Lancker, D., & Kreiman, J. (1987). Voice discrimination and recognition are separate abilities. Neuropsychologia, 25, 829-834.

Warrington , E.K., & Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. Neuropsychologia, 12, 419-428.

Whittlesea, B. W. A., & Cantwell, A. L. (1987). Enduring influence of the purpose of experiences - encoding-retrieval interactions In word and pseudoword perception. Memory & Cognition, 15, 465-472.

Zatorre, R.J., Evans, A.C., Meyer, E., & Gjedde, A. (1992). Lateralization of phonetic and pitch discrimination in speech processing. Science, 256, 846-849.

Ziemer, H., & Bowers, J. S. (1999). Phonological and orthographic contributions to long-term priming in the lexical decision task. Manuscript submitted for publication