VISUAL COGNITION, 1999, *6* (3/4), 461–492

Temporal Integration of Information in Orthographic Priming

Jonathan Grainger
CNRS and Université de Provence, Aix-en-Provence, France

Arthur M. Jacobs
Phillips-University of Marburg, Germany, and CNRS, Marseille, France

We present a theoretical framework for studying the integration of information over time in experimental paradigms involving the successive presentation of printed strings of letters with varying orthographic overlap (orthographic priming). This framework is expressed as a set of six processing principles, five of which are derived directly from the interactive activation model of orthographic processing (McClelland & Rumelhart, 1981), and an additional sixth principle relative to a hypothesized inhibitory reset mechanism. The validity of these principles are evaluated with respect to data obtained from different variants of the orthographic priming paradigm, and related paradigms.

TEMPORAL INTEGRATION AND DYNAMIC MODELS OF INFORMATION PROCESSING

Fluent reading requires rapid access to individual word meanings. This process is facilitated by the spaces that mark word boundaries in printed text. When the eyes fixate a given word during reading, information concerning letter identities and their respective positions must be rapidly translated into semantic information. In the case of monomorphemic words (i.e. words that do not contain affixes and that are not compounds), this translation process must operate via some form of whole-word representation. In the present paper, we examine how orthographic information accumulates over time, allowing a unique whole-word representation in memory to be isolated for further processing.

J. Grainger, CREPCO, Université de Provence, 29 Av. Robert Schuman, 13621 Aix-en-Provence, France. E-mail: grainger@ newsup.univ-mrs.fr

The authors thank Professor Glyn Humphreys for his detailed comments on an earlier version of this work.

One paradigm appears particularly appropriate for examining the integration of orthographic information over time during visual word recognition. This is the orthographic priming paradigm, which involves the successive presentation of strings of letters. It is hypothesized that subjects' performance in a given laboratory task will be affected by the way in which information extracted from these successive stimuli is integrated over time. As noted by Forster (1993), this integration process may operate at different levels. It is the careful manipulation of factors thought to affect the integration process at a particular level, such as the orthographic overlap between stimuli, and the precise timing of the prime-target sequence that should help uncover the underlying mechanisms.

In order to understand how successive events in time are integrated (or not) into coherent wholes, as well as an appropriate paradigm, an appropriate theoretical framework is required. This framework must be able to generate precise predictions concerning the temporal aspects of human information processing. However, in the field of visual word recognition, static rather than dynamic models of information processing are the rule. These can be verbal-boxological models (e.g. Forster, 1976), mathematical models (e.g. Massaro & Cohen, 1991), or "one-step" connectionist models (e.g. Seidenberg & McClelland, 1989).[1] Such models only allow predictions at the level of asymptotic performance, and as such ignore a wealth of useful experimental data. With respect to the present work, the major drawback of such static models of visual word recognition is that they do not allow precise predictions concerning the temporal evolution of orthographic priming effects.

The family of interactive activation models of human information processing is a good example of the application of non-linear dynamical systems in cognitive psychology, while retaining localist representations for tractability (see Grainger & Jacobs, 1998, for a discussion of the pros and cons of localist representations in connectionist models). In what follows we present an account of temporal integration of orthographic information during visual word recognition that is guided by the basic principles of the interactive activation family of models. In particular, our analysis will draw much inspiration from the prototypical interactive activation model of visual word recognition described by McClelland and Rumelhart (1981; Rumelhart & McClelland, 1982).

The notion of orthographic information refers to knowledge of spellings of words. In alphabetic languages such as English and French, we assume that such knowledge is essentially letter based. In other words, knowledge of how to

---

[1] Some static models may be superficially dynamic. Since a model must be judged (at least partly) by the predictions it can generate, only models that provide clear predictions concerning intermediate processing stages and the form of information accumulation functions are considered truly dynamic here (see Jacobs & Grainger, 1994, for further discussion).

spell a word is thought to be stored as a set of abstract representations that code both the identity and position of a word's component letters. It is also assumed that early visual processing operating on strings of letters is shared with all complex visual forms (e.g. objects, faces). This will involve the detection of alignments and discontinuities in luminance changes across the visual field by orientation sensitive cells and end-stopped cells in the visual cortex (Boucart, 1996). However, in the case of printed strings of letters, it is hypothesized that the output of early visual processing serves as input to a highly specialized horizontally oriented bank of "letter" detectors. This letter detector array serves to "bind" the appropriate featural information around letter identities at a specific position in the stimulus. These are hypothesized to be noisy letter detectors in that some amount of feature-letter, cross-channel leakage is assumed (e.g. by applying a gaussian spread to neighbouring positions, see Peressotti & Grainger, 1995). Nevertheless, such noise will be greatly reduced in the case of printed alphabetic stimuli compared to cursive (i.e. continuous handwritten) script.

It is further assumed (following McClelland & Rumelhart, 1981), that these letter detector arrays feedforward information to whole-word orthographic representations (see also Grainger & Jacobs, 1996). Information accumulates in these letter and word detectors over time until a critical amount of information has been extracted from the stimulus, thus enabling word recognition to be achieved. The interactive activation model or IAM (McClelland & Rumelhart, 1981) provides a detailed, transparent, algorithmic description of these processes. Non-linear activation functions govern feature-letter and letter-word information flow. Thus, activation builds up over time in the processing units corresponding to letters and words until a critical or asymptotic level has been reached. By presenting successive stimuli to this dynamic model one can examine (in this model world) how different sources of orthographic information are integrated over time. Such a model therefore makes precise predictions with respect to how effects of orthographic priming evolve over time. These predictions can then be compared to data obtained with human participants in orthographic priming experiments.

However, rather than presenting simulation work with the dynamic algorithmic model (see Jacobs & Grainger, 1992; Grainger & Jacobs, 1996, for examples of this), here we attempt to derive a set of verbalized processing principles (Jacobs & Grainger, 1994) from this specific model, that provide a higher-level conceptualization of the integration of orthographic information over time. Simulations are necessary to generate precise quantitative predictions from the model, in order to test it against human data. However, modelling primarily serves understanding, hence the present attempt to express what we have understood about orthographic processing from our travels in the model world of interactive activation. But, before continuing our theoretical analysis, first a methodological note on priming.

METHODOLOGICAL CONSIDERATIONS

The priming paradigm has been widely used to study the hypothesized component processes of visual word recognition (orthographic, phonological, morphological, semantic, and syntactic). In what follows we describe some recent developments in priming methodology that are particularly relevant to the question of temporal integration of information during visual word perception. These developments concern the use of very brief prime exposure durations and the precise manipulation of prime exposure and/or intensity.

Standard priming

It should be noted (in order to avoid further confusion), that the term "prime" is used here in its purely methodological sense without any bias in terms of its theoretical interpretation. Priming methodology offers a direct technique for examining the question of temporal integration in human information processing. Two or more separate events are presented to the observer in succession, and the resulting influence on behaviour is measured in an appropriate response-generating paradigm. Differences in performance (i.e. speed and/or accuracy of response) to a given target stimulus preceded by different prime stimuli are taken to reflect differences in the way information extracted from the prime stimulus has influenced subsequent target processing.

There are many ways that the presentation of stimulus 1 (the prime) can influence processing of stimulus 2 (the target), and only some of these are relevant to the present analysis. For example, when a hypothetical participant (typically a not so naïve student in psychology) sees a series of primes and targets that are strong semantic associates (e.g. DOCTOR-NURSE), then on presentation of the prime BREAD our participant may well be inclined to guess that the following word will be BUTTER. Evidence for such strategic "predictive priming" effects take the form of: (1) relatedness proportion effects; and (2) differences between unrelated and "neutral" conditions (see Neely, 1991 for a review). Concerning point 1, when one increases the percentage of related versus unrelated prime-target pairs in an experiment, semantic/associative priming effect sizes increase, but only at relatively long stimulus onset asynchronys (SOAs > 500 ms). Concerning point 2, one observes a deficit (slower RTs and more errors) when comparing unrelated primes (e.g. TREE-NURSE) to "neutral" primes (e.g. XXXX-NURSE), but once again only with relatively long SOAs (and not in the word-naming task). A neutral (XXXX) prime cannot be used to predict a target, whereas an unrelated prime (e.g. TREE) could generate a false prediction (e.g. LEAF when the target is NURSE). Although such predictive/strategic priming effects are of interest in themselves (e.g. Stone & Van Orden, 1993), the present work concentrates on fast, automatic priming effects that are not influenced by such predictive strategies.

Masked priming

The masked priming paradigm was developed initially to investigate the effects of subliminally presented prime stimuli (e.g. Marcel, 1983). To this end, prime stimuli are presented very briefly and preceded and followed by a pattern mask (although in many cases the target itself acts as a backward mask).[2] The interesting aspect of the masked priming paradigm, for the present purposes, is that subliminally presented primes cannot give rise to predictive strategies. Thus, any influence of prime presentation on performance to target stimuli can be assigned to information automatically extracted from the prime stimulus in the absence of any awareness of such processing (see Forster, 1993 for a discussion of the advantages of masked priming).[3]

The first published application of the masked prime paradigm to the study of purely formal prime-target relatedness, was the study by Naish (1980). This author studied the influence of phonological and visual (orthographic) prime-target overlap using the backward masking variant of the masked prime paradigm. In this situation, targets are very briefly preceded by a pattern mask and followed by the prime stimulus which acts as a backward mask for the target (see also Perfetti, Bell, & Delaney, 1988). The prime stimulus is itself briefly presented and followed by a pattern mask. Participants were requested to report both the target and the prime-mask. Naish (1980) demonstrated greater accuracy of target word report when prime-masks were visually similar (e.g. choir-chair) or homophones of the target (e.g. choir-quire), compared to an unrelated condition (e.g. choir-seat). Evett and Humphreys (1981) and Humphreys, Evett, and Taylor (1982) published similar work using a forward priming technique with the perceptual identification task. In their experiments both prime and target stimuli were briefly presented in succession, preceded and followed by a pattern mask (i.e., mask-prime-target-mask: The four-field priming technique). Targets were presented in upper case and primes in lower case letters, and participants were instructed to report the word presented in upper case. These experiments clearly demonstrated that primes sharing orthographic and phonological information with targets improved target recognition relative to an unrelated control condition. Later studies demonstrated that

---

[2] Some researchers have preferred the so-called backward masking paradigm (e.g. Naish, 1980) where one examines the influence of a post-target stimulus on identification of the target word. Although strictly speaking this is not a member of the family of priming paradigms (since primes logically precede targets), we will include it as such in order to simplify the following discussion.

[3] There is, of course, a long-standing debate with respect to whether or not prime stimuli are indeed unidentified by participants in so-called subliminal priming experiments (see e.g. Greenwald, Klinger, & Schuh, 1995). However, this debate has until now taken a purely methodological form (i.e. what is the best measure of prime identifiability). In the present work we discuss the consequences of prime stimuli becoming identifiable within our theoretical framework.

phonological prime-target similarity has an effect on its own when measured against appropriate orthographic controls (e.g. Perfetti & Bell, 1991; Perfetti et al., 1988). What remained to be demonstrated was whether or not prime-target orthographic overlap influenced target processing in the masked prime paradigm, when measured against phonological controls.

This specific question was addressed by Ferrand and Grainger (1994) and Grainger and Ferrand (1996) by using pseudohomophone primes (i.e. non-words that are pronounced like a real word, such as BRANE-BRAIN) that differ in terms of their orthographic overlap with the target. These primes are difficult to find in English but can be generated in French (e.g. NERF is a French word, "nerve", that is homophonic with the following two non-words: NERT and NAIR).[4] By comparing the relative influence of NERT and NAIR type primes on performance to target words, effects of prime-target orthographic overlap can be estimated in the absence of any confound from phonology (in both cases there is 100% phonological prime-target overlap). The results show that lexical decision RTs to target words are facilitated by increased orthographic overlap at prime exposure durations of 29 and 43 ms. With longer prime exposures, these facilitatory effects tend to disappear.

The studies of Ferrand and Grainger (1994) and Grainger and Ferrand (1996) used a variant of the masked prime paradigm introduced by Forster and Davis (1984). The critical difference with respect to the paradigm used by Evett and Humphreys (1981) and Humphreys et al. (1982) is that target duration is increased such that subjects can perform a lexical decision on the target at high levels of accuracy. The sequence of events in the Forster and Davis variant is typically a forward mask for approximately 500 ms followed immediately by the prime stimulus for about 60 ms and replaced by the target that remains on the screen either for a fixed duration (typically 500 ms) or until participants respond. In spite of the similarity in the perceptual identification and lexical decision variants of masked priming, Davis and Forster (1994) have noted one important difference in terms of their relative sensitivity to low-level feature summation. Although primes and targets are typically presented in different case, they are almost always presented at exactly the same screen position. This allows for a certain amount of purely physical overlap between lower and upper case stimuli that share letters in the same position. Davis and Forster demonstrated that masked priming with the perceptual identification task was sensitive to such low-level visual overlap in the absence of higher-level orthographic and/or phonological overlap. The lexical decision variant of

[4] The feasibility of this in French is due to the combination of relatively consistent spelling-to-sound mappings in French (at least at the level of rime units) plus the inconsistent mapping from sound to spelling. Thus, the letter sequence ERT can only be given one pronunciation /er/, but the sound /er/ can be assigned multiple written forms (ERT, ERF, ERC, AIR, ERE). See Ziegler, Jacobs, and Stone (1996) for an analysis of these correspondences in French.

masked priming did not, however, show such sensitivity. The experiments of Davis and Forster (1994) suggest that it is the extended duration of the target in the lexical decision task (rather than differences in prime exposure duration or type of response) that is the critical factor determining whether or not such low-level visual summation effects arise.

Incremental priming

Dynamic models of human information processing describe how information builds up and how different types of information interact over time. The incremental priming technique (Jacobs, Grainger, & Ferrand, 1995) was developed in order to provide appropriate data on this point, and to establish a new baseline for priming studies. The technique consists in a gradual increase of the prime's informational value (operationalized as prime intensity or duration). The minimum prime intensity/duration level serves as a within-condition baseline for each priming condition. This corresponds to some (usually very low) value of prime intensity and/or duration where preliminary experimentation has demonstrated no differences across priming conditions. Thus, we can define any priming effect with respect to two baseline conditions, one is the minimum intensity/duration condition of the particular priming condition (within-condition baseline), and the other is a different priming condition (across-condition baseline). This double-baseline approach makes measuring of priming effects more reliable and imposes stronger constraints on our interpretations of these effects.

This new baseline, the within-condition baseline, serves as an additional guard against premature conclusions concerning the direction of a priming effect (i.e. whether it is facilitatory or inhibitory). As argued by Jacobs et al. (1995), the standard way of interpreting effects obtained in priming experiments (i.e. an effect is defined as a difference between two conditions that use different types of prime), contains potential flaws, due to the arbitrariness of defining the neutral or baseline condition.

Consider the following examples. In an alphabetic decision task (speeded letter/non-letter classification) using masked primes, a blank character has been used as a "neutral baseline" (e.g. Jacobs & Grainger, 1991, Exp. 1). However, it can be argued that a blank prime gives the subject nothing to process and therefore is not directly comparable with other priming conditions, in which there is always something to process. However, choosing a "*" or something like it as the "neutral" prime is just as bad, because, given the letter/non-letter classification task, it can be argued that a "*" is not neutral in that it may bias the decision system toward a non-letter response (Arguin & Bub, 1995; Jacobs & Grainger, 1991). The same logic can be applied to a lexical decision task in which a row of "Xs" or "*s" are standard "neutral" primes, although they may very well create a No-response bias (e.g. Neely, 1977; for a critique see Antos, 1979; de Groot, 1984). If a row of “Xs” produces response competition in the LDT (see Jacobs et al., 1995), the many effects that have been interpreted as facilitatory in comparison with that presumably neutral condition in many classic studies should be very carefully reconsidered.

The incremental priming technique offers a potential remedy to this baseline or neutral condition dilemma. In the second baseline condition used in this technique, prime intensity (or duration) is set at a level sufficiently low to yield a null priming effect. The zero point is thus defined as the level of prime energy or availability for which no influence on target processing is observed. Then prime intensity is gradually increased. If RT decreases in comparison with the zero condition, this is taken as evidence that increasing the availability or information value of the prime *facilitated* target processing. If RT increases, we take this as evidence for *inhibition* of target processing. Of course, one can continue to use the terms “facilitation” and “inhibition” when discussing results from non-incremental priming studies, when referring to the fact that performance is better or worse in one priming condition compared to another. One must simply be more careful in the interpretation of such effects, being aware that between-condition facilitation may reflect reduced within-condition inhibition rather than true facilitation.

Jacobs et al. (1995) argue that by use of a double-baseline approach, the incremental priming technique makes measuring of priming effects more reliable and imposes stronger constraints on our interpretations of priming effects. On the one hand, the within-condition baseline cannot replace theoretically motivated across-condition comparisons that aim at determining the locus of priming effects. On the other hand, it is difficult to claim that a prime that yielded shorter RTs than in a given neutral baseline, but did not shorten RT with increases in its intensity and/or duration (incremental priming), really facilitated target processing. Thus, the double-baseline technique provides double-safety for theoretical interpretations of priming effects. It indeed puts stronger constraints on models of priming if one respects the following rule: Both significant *across-condition and within-condition priming effects* should be measured in theoretically crucial priming experiments and/or conditions. If only one of these two priming effects is obtained, the data are not clearly interpretable. An analogy with psychopharmacological methodology maybe helpful here. If, in a psychopharmacological study, one obtains a significant drug vs. placebo effect, but no dose effect (i.e. drug level 1 vs. level 2), the data are generally considered as not interpretable. Jacobs et al. (1995) simply propose a similar logic for priming studies.

Finally, the incremental priming technique provides critical data for testing dynamic models of human information processing that predict how priming effects vary as a function of prime exposure duration. This is particularly important when a given (within-condition) priming effect appears only within a very precise range of exposure durations (see e.g. Ferrand & Grainger, 1993).

Using a single prime exposure duration could lead to the premature conclusion that there is no effect of a given variable.

A THEORETICAL FRAMEWORK

We have recently presented a multiple read-out model (MROM) of orthographic processing in visual word recognition (Grainger & Jacobs, 1996) designed to explain variations in effects of orthographic variables across different experimental tasks. This model was developed within the general framework of an interactive activation model of visual word recognition (McClelland & Rumelhart, 1981). In the present analysis we attempt to isolate those aspects of the interactive activation framework that critically determine its account of orthographic priming effects. This principle-based modelling strategy, discussed at length by Jacobs and Grainger (1994) and Stone and Van Orden (1994), should help clarify the utility of such a theoretical approach for understanding results obtained with orthographic priming and related paradigms. Finally, we propose some key predictions generated by this approach to be tested in future experimentation.

Outline of the theoretical framework

In what follows, we analyse and extend an existing model of orthographic processing (the interactive activation model) in an attempt to evaluate its explanatory adequacy with respect to orthographic priming phenomena. In so doing we apply the strategy of "nested modelling" (Grainger & Jacobs, 1996; Jacobs & Grainger, 1994) whereby a new model should include the old model as a special case, or dismiss with it after falsification of the core assumptions of the old model. Our analysis, inspired by Marr's (1982) analysis of information processing systems, is divided into two components: The representational framework (this defines the different representations occurring at different levels of processing in a hierarchical information processing system), and the processing principles that critically determine the performance of the information processing device.

Table 1 summarizes the representational framework derived from the interactive activation model of visual word recognition. The following minimal set of principles are an attempt to verbalize the computations that can operate on the representations defined in Table 1.

Principle 1: Hierarchical, cascaded, non-linear, activation flow. On presentation of a string of letters, activation builds up over time in letter representations and continuously feeds forward to O-word representations.

TABLE 1
A Simplified Representational Framework for the Interactive Activation Model of Printed Word Perception (McClelland & Rumelhart, 1981) following Marr (1982).

Level	*Purpose*	*Primitives*
Letters	Orthographic representation of the stimulus	Letter-centred featural information
O-words	Orthographic representation of known words	Word-centered letter identities

Principle 2: Letter-word activation rule. Word unit activation strength is a function of the degree of orthographic overlap[5] with the stimulus. Activated letter units excite compatible O-words and inhibit incompatible O-words.

Principle 3: Top-down feedback (resonance). All positively activated O-words send positive activation to their component letter representations.

Principle 4: Lexical inhibition. All positively activated O-words send inhibition to all other O-words activated above zero.

Principle 5: Resting level activation. O-Word representations have resting level activations proportional to their (log transformed) frequency of occurrence.[6]

Applications to orthographic priming

According to the representational framework (Table 1) and the processing principles described previously, what critically determines the size and the direction of orthographic priming effects is the activation state of letter and word representations at target onset: (1) Letter unit activation is determined by the presence or absence of letters in the prime stimulus (P1), word-letter feedback (P3) and prime intensity and/or duration (P1); (2) Word unit activation is determined by the degree of orthographic overlap between word representations and the prime stimulus (P1, P2), the printed frequencies of all positively activated

[5] Degree of prime-target orthographic overlap clearly depends on how the position of each letter in a stimulus string is coded and how such information is stored in memory. In the present paper we follow McClelland and Rumelhart (1981) in applying the simplifying assumption that letters are coded for their specific position as a function of string length (e.g. second letter in a six-letter string).

[6] In Morton's (1969) logogen model, this if formally equivalent to having response/identification thresholds for individual words vary as a function of their printed frequency. One must note, however, that the formal equivalence no longer holds for the IAM due to within-level and between-level dynamics.

word units (P5), mutual inhibition between all positively activated word units (P4), and prime intensity and/or duration (P1).

The independent variables that will affect the size and direction of orthographic priming are therefore: (1) Degree of prime-target orthographic overlap– the number of letters shared by prime and target and their relative position. (2) Prime lexicality– the lexical status (word or non-word) of the prime and its printed frequency if it is a word. (3) The orthographic neighbourhood shared by primes and targets– the number of words that have maximum orthographic overlap with both the prime and the target. (4) Prime duration/intensity.

EVALUATION OF THE THEORETICAL FRAMEWORK

The theoretical framework provides clear qualitative predictions concerning how various independent variables may influence the effects of orthographic priming. Some of these predictions have already found support in currently available data, others still await testing in future experimentation. Here we examine evidence for effects of the four independent variables thought to critically determine orthographic priming effects as derived from the above analysis.

Degree of prime-target overlap

The vast majority of studies investigating effects of orthographic priming have applied Coltheart, Davelaar, Jonasson, and Besner's (1977) metric of orthographic similarity to define prime-target orthographic overlap. A word that has maximum orthographic similarity (without identity) with another word is called an orthographic neighbour of that word. These orthographic neighbours share all but one letter, respecting word length and letter-in-word position (e.g. WORD, WORK).

Studies that actually vary the degree of prime-target orthographic overlap (rather than contrasting the extreme cases of neighbour primes vs. unrelated prime) are quite rare. One example is the work of Humphreys, Evett, and Quinlan (1990). Using the four-field masking procedure where both prime and target are briefly presented in succession, preceded and followed by a pattern mask, these authors observed a monotonic increase in the size of orthographic priming effects as a function of the degree of prime-target orthographic overlap. Primes are presented in lower case and targets in upper case. Participants are asked to report the word presented in upper case. However, as previously noted, Davis and Forster (1994) have pointed out a potential problem with this particular version of the masked prime paradigm. These authors noticed that variations in the legibility of superimposed prime and target stimuli (with unlimited viewing time) influences performance to target stimuli in the four-field masking paradigm. This legibility effect does not arise in the lexical decision variant of the masked prime paradigm. With respect to this point, it is therefore important to note that Humphreys et al.'s work has been recently replicated using response time (RT) in the lexical decision task as the dependent measure (Peressotti & Grainger, in press). Moreover, Peressotti and Grainger used different sized prime and target stimuli in order to avoid any confound with low-level physical summation. Finally, the robust relative-position priming effects (e.g. BVK-BLACK) obtained by Humphreys et al. in the perceptual identification task, are unlikely to be due to a confound with superimposed prime-target legibility.

One interesting aspect of the recent work by Peressotti and Grainger is that net orthographic priming effects appear to reflect the combination of facilitation from shared letters and inhibition from different letters. This trade-off was suspected to be the origin of the difficulty in observing any significant facilitatory effect of SDDS primes (where S represents a letter shared with the target and D a different letter) compared to DDDD primes. In an incremental priming study (see earlier section), S++S primes (where different letters were replaced by a + sign) were shown to produce both within-condition and between-condition facilitation, while SDDS primes gave no significant within-condition priming. The all-different letter primes (DDDD) produced within-condition inhibition (i.e. longer RTs to target words with increasing prime exposure duration). These inhibitory effects of different letters are explained as a direct consequence of incompatible letter-word inhibition (principle 2) in the interactive activation model.

Prime lexicality

A quick survey of all orthographic priming studies using the masked prime procedure shows that the lexical status of the prime critically determines the direction of priming effects relative to an unrelated baseline. When primes are orthographically related words (e.g. blink-BLANK) presented for approximately 60 ms, then responses to the target word are typically slower compared to an unrelated word prime (e.g. trout-BLANK). This between-condition inhibition is strongest when the prime is a high-frequency word and the target low frequency (Segui & Grainger, 1990). On the other hand, primes that are orthographically related non-words (e.g. blonk-BLANK) tend to facilitate responses to the target word compared to unrelated prime stimuli (Ferrand & Grainger, 1992, 1993, 1994; Forster, 1987; Forster, Davis, Schoknecht, & Carter, 1987; Perfetti & Bell, 1991; Perfetti et al., 1988). Note that a similar interaction between prime lexicality and the size and direction of priming effects has been reported in unmasked auditory priming studies (e.g. Slowiaczek & Hamburger, 1992).

Thus, all other things being equal, orthographically related word primes tend to interfere with target processing whereas orthographically related non-word primes tend to facilitate target processing relative to unrelated controls. However, prime duration has been shown to influence these effects. As one increases prime exposure duration up to approximately 80 ms then the inhibitory effects of orthographically related word primes steadily increase (Grainger, 1992). With respect to non-word primes, Ferrand and Grainger (1994) have shown that, measured against a phonological control, orthographic priming effects are facilitatory with prime exposures around 30 ms and disappear as prime exposures reach 50–60 ms.

In the theoretical framework described previously, the application of principles 1, 2, and 4 provides an explanation of this pattern of facilitatory and inhibitory orthographic priming effects. Activation builds up in letter representations before word representations. Prime-target orthographic overlap produces initial facilitation from letter unit activation (principles 1 and 2). However, to the extent that prime and target stimuli support the activation of words other than the target, within-level inhibition on the target word will develop (principle 4), and this will be maximal when primes are real words. This within-level inhibition may then cancel the bottom-up facilitation generated by shared letter representations (this will depend on the shared neighbourhood of the prime and target– to be discussed later).

Shared orthographic neighbourhood of the prime and target

The first relevant data with respect to this issue were reported by Forster and his colleagues (Forster, 1987; Forster et al., 1987). Forster found that in the case of non-word primes that were orthographic neighbours of target words, facilitatory effects relative to an unrelated baseline were not robust with short words (four letters long). Further work indicated that word length *per se* was not the critical variable, but had been confounded with the neighbourhood density of the prime/target (short words have more orthographic neighbours than large words when applying Coltheart et al.'s, 1977 definition). Further experimentation demonstrated that short words with few orthographic neighbours produced significant priming. Later work by Forster and Davis (1991) allowed these authors to confirm that it is indeed the neighbourhood density of the target word that critically determines the size of priming effects in the masked prime paradigm. Facilitatory priming arises only in targets with small numbers of orthographic neighbours (i.e. targets with low neighbourhood density).

However, within the theoretical framework developed, target word neighbourhood density *per se* should not influence orthographic priming. The critical factor is the degree to which the target word continues to provide bottom-up support to words (other than the target itself) that were activated by the prime stimulus. This is determined not by the overall number of orthographic neighbours of the target but by the number of orthographic neighbours of the target word that are also orthographic neighbours of the prime stimulus. However, target word neighbourhood density, as manipulated by Forster and his colleagues, is confounded with the orthographic neighbourhood shared by prime and target. High-density targets will tend to share many orthographic neighbours with the prime (e.g. CEAL and DEAL have heal, real, meal, peal, etc. as shared neighbours), whereas low-density targets will tend to be the only orthographic neighbour of the prime (e.g. VINC-ZINC). According to our account of orthographic priming, it is the shared neighbourhood of prime and target that is critical. Facilitatory effects should be strongest when primes and targets have no other neighbours in common.

Thus, one key prediction of the model of orthographic priming presented here, is that the shared neighbourhood (SN) of the prime and target rather than the neighbourhood size of the target itself, should determine the size of priming effects. One recent study (Hinton, Liversedge, & Underwood, 1996) has found evidence that this is indeed the case. Hinton et al. (1996) manipulated the shared neighbourhood of primes and targets in an orthographic priming experiment. These authors varied what they termed the "ambiguity" of primes defined in terms of the number of orthographic neighbours that could be generated from the prime. Thus, the target word CODE was primed by an ambiguous prime _ODE (compatible with code, mode, rode, bode, etc.) and an unambiguous prime C_DE (only compatible with code). Measured relative to an unrelated condition, only unambiguous primes gave significant facilitation at 100 ms prime exposures.

A similar result was found in a recent unpublished study.[7] The orthographic neighbourhood shared by prime and target was manipulated in high-density target words (i.e. those where no facilitatory priming should be obtained according to the density constraint hypothesis, Forster & Davis, 1991). Either the target word was the only word in the prime-target neighbourhood (SN–), or the shared neighbourhood contained several other words (SN+). The same target word was primed with a SN– prime, a SN+ prime, and an unrelated prime. The prime presentation duration was either 14, 29, 43, or 57 ms. This study replicated the statistically fragile nature of facilitatory form priming in high-density words at the longest prime exposures (i.e. those used typically in the studies of Forster and colleagues). However, this was only true for the SN+ primes relative to the unrelated prime condition. The SN– primes produced robust (by participant and by item) facilitation relative to the unrelated prime condition at the longest prime exposures. As the exposure duration of the prime

[7] This unpublished work was conducted in collaboration with Elsa Spinelli in Paris, and Ton Dijkstra and Walter van Heuven in Nijmegen.

increased, RTs associated with SN− primes remained stable, whereas RTs to the SN+ prime condition increased. Our interpretation is that all words that are neighbours of both the prime and the target receive maximum bottom-up input (principle 2) and therefore generate the strongest within-level inhibition on the target word (principle 4). This lexical inhibition develops with increasing prime exposure. These results, along with those observed by Hinton et al. (1996), therefore suggest that the shared neighbourhood of primes and targets does influence masked orthographic priming.

Below and above threshold priming

In the theoretical framework presented in the previous section, with prolonged exposure to a given stimulus word, activation rises in letter and O-word representations to asymptotic values. When the stimulus is removed, these representations then slowly decay back to their resting-level activations. However, the model lacks an efficient means of dealing with the rapid succession of printed words that occurs in normal reading. In its present form, activation from the directly preceding word will cause massive inhibition during the processing of a different word that follows. This situation has been remedied in other interactive activation models (e.g. Dell, 1986) by applying what is often referred to as a "reset" or "activation dumping" mechanism (Meiran, 1996), or "self-inhibition" (MacKay, 1987). There are several means of implementing such a mechanism in an interactive activation framework. One commonly used mechanism in models of production (Dell, 1986; Rumelhart & Norman, 1982) is an immediate return to resting-level activation of representations that have been selected for output. In models of perception, on the other hand, one can use a mismatch between incoming sensory information and feedback from activated representations to trigger an inhibitory reset. Such a mismatch-triggered inhibitory process can be found in adaptive resonance models of perception and categorization (e.g. Grossberg, 1987). In the following, we provide a tentative verbal description of such a mechanism (expressed as an additional sixth principle in the interactive activation framework), and we examine how it can account for different patterns of above-threshold orthographic priming.

Principle 6: Stimulus changes and inhibitory reset. Stimulus changes do not interrupt processing when the change occurs within the limits of a critical time/intensity limit of the first stimulus. After this critical moment, a *reset mechanism* causes the momentary inaccessibility of all activated representations when a "new" stimulus is detected. It is further hypothesized that *the amount of inhibition directed toward a given representation is proportional to the activation level of that representation* when the reset mechanism comes into operation.

In a specific implementation of this reset mechanism, letter-level activity generated by bottom-up information could be compared with letter-level activity generated by top-down information. One would, of course, require two sets of "letter-level nodes" for such a mechanism to operate (e.g. the STM and LTM nodes in adaptive resonance theory– (ART); Grossberg, 1987). When the mismatch between these two sources of information exceeds a certain criterion value, then inhibition is sent to all activated word-level representations. The mismatch criterion therefore determines whether or not two successively presented strings of letters are integrated or not into a single perceptual event.

A critical duration/intensity of the first stimulus is necessary before triggering a reset, such that sufficient higher level semantic information has been extracted before lower level orthographic codes are inhibited. We further assume that when the second stimulus is a blank (or some other non-linguistic stimulus that does not resemble a letter string), then the representations activated by the first stimulus simply decay in the absence of further stimulation and in the absence of a mismatch. In other words, the degree of mismatch is zero if the second stimulus does not significantly activate any letter representations. Thus, there are three necessary conditions for an inhibitory reset to occur when two stimuli are presented in succession: (1) one-word representation must have attained a sufficiently high enough activation level during the processing of the first stimulus; (2) the second stimulus must provide minimal activation input to letter representations; and (3) the mismatch between the two stimuli must reach a criterion level. In what follows, we provide an account of orthographic priming effects obtained with unmasked visible primes based on the operation of an inhibitory reset mechanism.

Colombo (1986) was one of the first to report inhibitory effects of unmasked orthographically related word primes in the visual lexical decision task. In Colombo's (1986) experiments, primes were Italian words or non-words that either shared their orthographic and phonological rime with target words (e.g. palmo-CALMO), or were totally unrelated to the target. The prime-target SOA was either 240 or 640 ms. In these conditions the related non-word primes facilitated target recognition relative to unrelated non-word primes, whereas related word primes inhibited target recognition. The inhibitory effects were strongest for high frequency target words.

Pursuing the work of Colombo (1986), Segui and Grainger (1990) tested the effects of orthographically related prime stimuli in conditions where primes were clearly visible and identifiable (350 ms prime exposures). Compared to unrelated primes, significant inhibition was observed only when primes had low printed frequencies and targets were high frequency. Similarly, Grainger (1990) observed inhibitory effects of low- and medium-frequency orthographically related primes on performance to high-frequency targets in the lexical-decision and word-naming tasks. Lupker and Colombo (1994) provided further evidence for inhibitory form priming effects in the lexical decision task only for high-frequency targets. These inhibitory effects were apparent at 140 ms and 315 ms SOAs, but turned to facilitation at a longer 805 ms SOA. Thus, it would appear from this set of results (Colombo, 1986; Grainger, 1990; Lupker & Colombo, 1994; Segui & Grainger, 1990) that both target word frequency and SOA determine whether or not inhibitory orthographic priming (measured relative to an unrelated condition) will be observed. As discussed previously, the effect of target frequency is explained by the higher activation level attained by target word representations during prime processing in the high frequency target condition. This higher activation level results in stronger inhibition during reset (principle 6). The use of long SOAs encourages subjects to use predictive strategies, leading to facilitatory priming effects (Lupker & Colombo, 1994).

Inhibitory effects of unmasked word primes have also been reported in auditory studies. When primes have high phonological similarity with target words (e.g. stiff-still, differing only by their final phoneme) then shadowing latencies (only name the target word) are slowed relative to conditions where primes and targets have less phonetic overlap (Hamburger & Slowiaczek, 1996; Slowiaczek & Hamburger, 1992). Slowiaczek and Hamburger (1992) systematically varied the degree of left-to-right phonetic overlap between prime and target words (e.g. target word STILL primed by still, stiff, steep, smoke, and dream). Low levels of similarity (one or two phonemes) produced facilitation relative to the unrelated condition, whereas maximum similarity (without identity) produced inhibition relative to the minimum overlap conditions. Hamburger and Slowiaczek (1996) demonstrated that the high similarity interference effect is not modifiable by expectancy manipulations (varying the proportion of phonologically related prime-target pairs and the interval separating primes and targets), whereas the low similarity facilitation effect is. They therefore conclude that the inhibitory effects are the result of automatic processes in auditory word recognition, such as the lexical inhibition mechanism (principle 5) proposed by McClelland and Rumelhart (1981) and further described in the present work (but see Dell & O'Seaghdha, 1994 for an alternative view).

Goldinger, Luce, and Pisoni (1989) in an auditory perceptual identification task (identify target words presented in noise) found inhibitory phonological priming effects that were only robust with low frequency primes. Goldinger et al.'s results suggest that it is the frequency of the prime rather than the target that critically determines the presence of inhibition. Nevertheless, in their Experiment 1A, it was the combination of a low frequency prime and a high frequency target that gave the strongest inhibition. This is the condition that gave the strongest inhibition in the unmasked studies of Segui and Grainger (1990). Indeed, this is the precise condition that will generate maximum target word activation during prime processing (excluding the case of prime-target identity), and therefore maximum inhibition of the target representation following prime recognition. Moreover, these inhibitory phonological priming effects disappeared when targets were presented 500 ms after prime offset. This may define the limits of effective inhibition following an inhibitory reset.

EXTENSION OF THE THEORETICAL FRAMEWORK TO RELATED PHENOMENA

The inhibitory reset mechanism discussed in the preceding section is a proposal shared (most often implicitly) by many parallel activation models of perception and production (see Meiran, 1996 for a recent discussion). In this section we will discuss how this particular mechanism can account for the empirical data on repetition priming and related phenomena. The term "repetition priming" is typically reserved (at least in the field of memory research) to refer to the facilitatory effects of repeating the same word (or other stimulus), with relatively long intervals (ranging from several minutes to one year!) between each occurrence of the stimulus (see Tenpenny, 1995 for a review). In what follows we examine the effects of repetition priming in situations with varying prime exposure duration and varying prime-target lag.

Principle 6 allows us to distinguish different forms of repetition priming as a function of whether or not an inhibitory reset is triggered, and whether or not the target is presented within the period of effective inhibition. As previously noted, one of the necessary conditions for an inhibitory reset to occur is that one-word representation reach a sufficiently high enough activation level during prime processing. When primes are exposed for short durations, no inhibitory reset is triggered and facilitatory repetition priming effects are observed (both within and between condition effects). These are the typical conditions tested in masked repetition priming to be discussed next, below. With brief prime exposures, perceptual integration of prime and target is assumed to occur. The interactive activation model provides one theoretical framework for describing such integrative processes. With longer (above-threshold) prime exposures, on the other hand, *type of intervening stimulus* (blank, pattern mask, or word of varying duration) and *target duration* critically determine the direction of prime-target repetition effects. These determine whether or not an inhibitory reset is triggered.

Masked repetition priming

Within our theoretical framework, masked repetition priming effects can only be facilitatory, whether measured across conditions (i.e. against an unrelated control) or within conditions (i.e. measured against a minimum intensity/duration baseline; Jacobs et al., 1995). Several studies have reported strong facilitatory effects of masked repetition priming measured against unrelated control primes (e.g. Forster & Davis, 1984; Grainger & Jacobs, 1993; Segui &

Grainger, 1990). More recently, Jacobs et al. (1995) and Giraudo and Grainger (in preparation) have examined the evolution of these masked repetition effects with increases in prime intensity or duration, using the lexical-decision task. As prime intensity or duration increases, the repetition priming effect increases in size, both with respect to a within-condition baseline (minimum intensity level) and a between-condition (unrelated word) baseline. In Jacobs et al.'s study, these repetition priming effects were of approximately the same size for primes and targets in the same or different case, thus suggesting that they do indeed reflect processes operating on abstract word representations. The net between-condition priming effects (RT in the unrelated condition minus RT in the repeat condition) observed at each prime exposure in Giraudo and Grainger's study are shown in Figure 1.

A linear regression of net priming effects (y) on prime exposure duration (x) gave the following equation: $y = 1.09x + 6.4$ (milliseconds); $r^2 = 0.78$. This demonstrates that prime and target are almost perfectly integrated in the masked prime paradigm. The size of priming effects are almost identical to the exposure duration of the prime. It should be noted that, as is usually the case, primes were presented in lower case and targets in upper case in this study. On this point, it is interesting to note that neither Grainger and Jacobs (1993) nor Jacobs et al. (1995) observed an influence of prime-target case congruity on masked repetition priming effects.

Ferrand (1996) examined masked repetition priming effects with the word naming task and variable prime-target intervals (ISI). Primes were presented for 29 ms followed by a backward mask for 14 ms, a variable delay (from 0 ms to 1000 ms), and then the target. Repetition priming was measured against a control word sharing the same initial phoneme as the target (e.g. table-TABLE

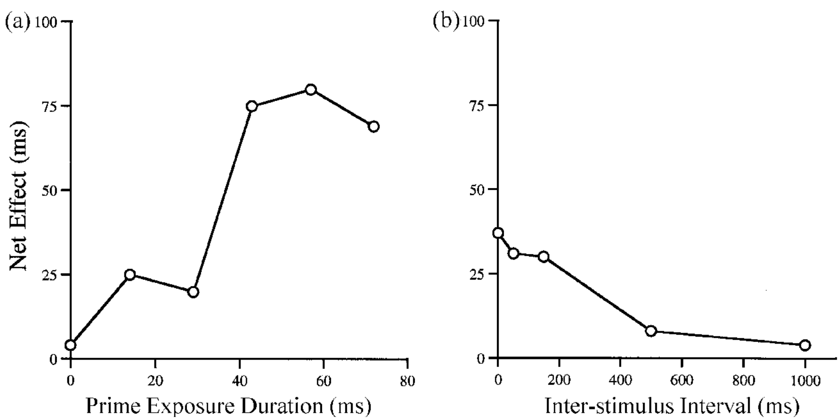

FIG. 1. (a) Time-course of information accumulation in whole-word representations as estimated by masked incremental repetition priming (Giraudo & Grainger, in preparation). (b) Time-course of decay of activation of whole-word representations as estimated by the variable SOA study of Ferrand (1996).

compared to tirer-TABLE) in order to control for onset priming effects that are prominent in the word naming task (see e.g. Grainger & Ferrand, 1996). These results are summarized in Figure 1 showing net between-condition effects of repetition priming (averaged across word frequency) as a function of ISI. Similar results showing rapid decay of masked repetition priming have been reported by Humphreys, Besner, and Quinlan, (1988).

Unmasked immediate repetition priming

Principle 6 allows us to generate predictions with respect to repetition priming effects involving identified primes and no different items intervening between prime and target. Immediate repetition priming is the condition where likelihood of an inhibitory reset is smallest. Repetition priming is the situation, by definition, where there is the least mismatch between prime and target, and therefore the lowest probability of triggering an inhibitory reset. This leads us to make the (intuitively very obvious) prediction that immediate repetition priming, contrary to orthographic priming, should produce strong facilitation.

Humphreys et al. (1988) have addressed this issue using the perceptual identification paradigm. The results of Humphreys et al. show that whether or not prime and target stimuli are separated by a pattern mask critically determines the direction of across-condition repetition priming effects. With 300 ms prime exposures followed immediately by target presentation, the repetition condition produced significantly lower levels of target identification compared to an unrelated prime condition. On the other hand, when primes were presented for 180 ms followed by a pattern mask for 120 ms (SOA = 300 ms) and then the target, repetition priming effects were facilitatory compared to the unrelated condition. Humphreys et al. argue that the intervening mask allows the prime and target to be perceived as separate perceptual events. They reported that prime stimuli were always correctly perceived independently of the presence of an intervening mask. In the absence of a mask, subjects apparently were perceiving the target stimulus as the continuation of the prime stimulus, and hence failed to report the target stimulus. In other words, the inhibitory effects of repetition priming observed with brief target exposures and in the absence of an intervening mask do not reflect the operation of an inhibitory reset. Rather, they reflect the perceptual integration of prime and target stimuli.

However, subsequent work on immediate repetition priming by Hochaus and Marohn (1991) and Hochaus and Johnston (1996) has complicated the picture somewhat. In the Hochaus and Marohn study, targets were presented immediately after prime stimuli but at a different spatial location (typically two lines below the prime stimulus). Participants had to identify the briefly (16 ms or 32 ms) presented targets (and also to judge target duration, which we will not discuss further here). When primes were presented for 250 ms (i.e. durations similar to those used by Humphreys et al., 1988), then the repeat condition gave superior target identification compared to the unrelated condition. This therefore replicates Humphreys et al., facilitatory repetition priming in conditions where prime-target perceptual integration is hindered. However, at 500 ms exposures, the repeat condition produced significantly lower target identification scores than the unrelated condition. This inhibitory effect of repetition inhibition was replicated in later work (Hochaus & Johnston, 1996) in a same-different judgement task (participants judged whether targets were the same as a probe-word presented 500 ms after the target). In these experiments prime words (referred to as "precues" by Hochaus and Johnston) were exposed for 250 ms, followed by a 115 ms mask, followed by targets exposed for 33 or 50 ms, finally followed by a mask, all in the same screen location. These experimental conditions (apart from the task requirements) were therefore very similar to the Humphreys et al. experiment with intervening masks. Yet, "same" judgement scores[8] were more accurate in the unrelated prime condition than the repeat condition in Hochaus and Johnston's experiments (e.g. 91% vs. 73% in Exp. 2). Thus, in very similar experimental conditions, participants find it easier to identify a briefly presented target (Humphreys et al., 1988), but find it harder to judge that the target is the same as a post-target probe word (Hochaus & Johnston, 1996), when the prime is the same as the target compared to an unrelated prime word. Clearly, more experimentation is in order here.

With longer target exposures (typically response limited), and no linguistic stimuli intervening between prime and target, then repetition facilitation should always be observed. This corresponds to the "zero lag" (where lag refers to the number of items intervening between prime and target) condition of what we will refer to as long-lag unmasked repetition priming, to be discussed next. Typically, the zero lag condition gives the strongest effects of repetition priming in these experiments (e.g. Monsell, 1985; Ratcliff, Hockley, & McKoon, 1985).

Long-lag unmasked repetition priming

Parsimony is a key goal for theory development in science. Critical advances can be made when apparently different phenomena are shown to have a common underlying cause. With this in mind, it is important to examine how theoretically motivated extensions of the interactive activation model can explain phenomena for which the original model was silent. Earlier we presented one such extension (principle 6). Here we argue for the importance of exploring the

---

[8] From our reading of the Hochaus and Johnston (1996) paper it appears that only "same" judgements were tested in both the repeat and unrelated prime conditions. "Different" judgements were never tested in the repeat condition (see their Table 1). We are assuming that participants were requested to make their same–different judgement only with respect to targets and post-target probes.

explanatory range of the basic model before appealing to such extensions. Principle 5 concerned modifications in resting-level activation of whole-word orthographic representations as a function of their frequency of occurrence. This principle is therefore tailor-made for explaining effects of printed frequency in visual word recognition. In what follows, we examine the extent to which this principle can be usefully invoked to accommodate data obtained with the long-lag repetition priming paradigm. This paradigm involves identified primes, long prime-target lags (in the order of minutes or longer) and items intervening between prime and target.

Following Morton (1969) and Monsell (1991), we assume that repetition priming, with long lags between the first and the repeated presentation of stimuli, is mediated by the same mechanism that mediates effects of printed word frequency. In Morton's logogen model, each identification of a given word resulted in a lowering of its response threshold over a considerable period of time. This lowered response threshold explained why the subsequent presentation of the same word many minutes or hours later, produced facilitation relative to an item that had not been previously presented. The systematic lowering of a logogen's threshold over subsequent repetitions gives rise to long-lasting modifications of the threshold, thus explaining the pervasive effects of word frequency. Exactly the same argument can be made concerning resting-level activations of whole-word orthographic representations in the interactive activation model (principle 5). Each presentation of a given word results in a rise in the resting-level activation of the corresponding whole-word representation. This increased resting-level activation slowly decays back to the original value unless subsequent presentations of the same word produce durable changes. It is these durable changes that are the basis of word frequency effects.

Certain results at present would appear to contradict this "abstractionist" (as opposed to episodic) account of long-lag repetition priming (see Monsell, 1985 for an early defence of the abstractionist position). Tenpenny (1995) summarizes these problems. First, with respect to the long-lag effects of repetition priming, Tenpenny claims that allowing word representations to remain in an "enhanced" state of accessibility over long time intervals should literally create havoc in the word recognition system (too many active words). However, it suffices to run simulations on the interactive activation model to show that this is not the case. Generally speaking, only words that share significant amounts of orthography with the target word are ever activated above zero, independently of resting level activation. Variations in resting-level activation mainly affect the activation levels attained by orthographic neighbours of the target word. Second, the fact that non-words show long-lag repetition priming (e.g. Feustel, Shiffrin, & Salasoo, 1983; Rueckl, 1990) has often been interpreted as a major obstacle for abstractionist accounts. However, in the theoretical framework presented here, orthographically legal pseudowords can produce significant levels of activation in word units, and therefore produce modified resting-level activations in these orthographically similar word units. During test, the higher activation levels of orthographically similar word units feeds back activation to letter units (principle 3), therefore improving pseudoword identification.

This analysis of long-lag repetition priming also explains why words presented at test that are orthographic neighbours of words presented during study will be better identified in the test phase (Rueckl, 1990). However, according to our theoretical framework, the influence of presenting orthographic neighbours during study will depend on the relative frequency of occurrence of the test word and its orthographic neighbour. Since the word presented during study will act as a competitor during the identification of the test word, and this competition will be exaggerated by its increased resting-level activation, the overall effect on test word identification will result from a trade-off between the benefits of a higher initial activation level and the cost of a stronger competitor. Unfortunately, only simulations run on an implemented version of this model will provide accurate predictions with respect to long-lag priming across orthographic neighbours, and the influence of their relative printed frequency. Moreover, our theoretical framework does not address issues of learning, and therefore ignores the possibility that pseudoword stimuli may (even after a single presentation) create a weak orthographic representation in long-term memory that would influence long-lag priming effects with these stimuli. This immediate learning could arise by associating a new word unit to simultaneously activated letter units (see Burton, 1994 for a discussion of Hebbian learning of new faces in an interactive activation model).

In strong favour of the abstractionist position, on the other hand, is the fact that morphologically related words (e.g. RAN-RUN) produce long-lag facilitation whereas formally related words (e.g. RUM-RUN) do not (e.g. Napps, 1989; Napps & Fowler, 1987). Furthermore, a critical test of the episodic account of long-lag repetition priming was recently provided by Dean and Young (1996). The rationale of their experiments is simple. If two stimuli (a word and a picture) are presented together at test, then the episodic memory of that event will retain information about the co-occurrence of the two stimuli. This predicts that stronger priming should obtain when the two stimuli co-occur again at test, compared to re-combinations of the stimuli (both of which were present during study but paired with different stimuli). The results show that this is not the case. Same-different judgement latencies were just as fast in the re-paired condition as the same pairing condition at the test phase.

It is therefore our opinion that an abstractionist account of long-lag repetition priming, as embodied in the present theoretical framework, provides a plausible account of the empirical results at present. One major advantage of this account is that it provides a unitary explanatory framework for all forms of priming. However, we quickly note that it is probably wise to leave open the possibility that repetition priming effects obtained with very long lags (i.e. hours or days rather than minutes) may be influenced by additional mechanisms. Only research systematically comparing the effects of the same variable across different delays will provide light on this question.

Repetition blindness

The reset mechanism presented earlier (principle 6) was used to account for "above-threshold" orthographic priming effects. In what follows we examine the extent to which this mechanism can provide a principled account of related phenomena. According to principle 6, when the necessary conditions for triggering a reset are met, there follows a short period of time during which stimuli corresponding to previously activated representations are harder to recognize. This phenomenon has been reported many times in the recent experimental literature, and has been referred to as "repetition blindness" or RB (e.g. Bavelier & Potter, 1992; Kanwisher, 1987, 1991; Kanwisher & Potter, 1990). There are many similarities between the paradigm used to investigate RB and standard immediate repetition priming discussed earlier. The temporal lag between repetitions that cannot be surpassed for RB to occur is in the order of 500 ms (Kanwisher & Potter, 1990; Park & Kanwisher, 1994).

RB does not occur across semantically related words (e.g. sofa-couch: Kanwisher & Potter, 1990), but does obtain across modality for items that refer to the same object (digit-word, picture-word: Bavelier, 1994; Bavelier & Potter, 1992). It is also observed across homonyms (e.g. she *rose* to pick the *rose* in the garden) and across orthographically and/or phonologically related words (e.g. car-cart; juice-loose: Bavelier & Potter, 1992; Bavelier, Prasada, & Segui, 1994; Chialant & Caramazza, 1997). Whether or not repetition of non-cognate translation equivalents (i.e. translates without form overlap) gives rise to improved recall (Altarriba & Soltano, 1996) or diminished recall (MacKay & Miller, 1994) remains an open question. MacKay, Abrams, Pedroza, and Miller (1996) suggest that RB does occur with translation equivalents when these are embedded in mixed-language sentences. In the present analysis we will concentrate on RB obtained with identical stimuli or with similar physical realizations of different semantic entities (MacKay et al.'s "surface blindness"), and ignore the problematical issue of whether or not RB obtains with sufficiently strong semantic overlap in the absence of form overlap.

The standard rapid serial visual presentation (RSVP) procedure used to investigate RB typically involves the succession of a series of items all presented at the same screen location[9] for durations around 100 ms per item. Repeated items generally do not occur first in the sequence, and are generally separated by at least one intervening item. Contrary to standard priming studies, however, in a typical RB experiment participants have to report as many

[9] Kanwisher (1991), and more recently Luo and Caramazza (1996), have reported RB using spatially distributed simultaneous presentation of items.

items as possible from the sequence (full report instructions). The unfortunate aspect of this procedure is that it is not easy to disentangle perceptual effects from memory effects (i.e. those due to processes specifically involved in retaining information for later report; but see Park & Kanwisher, 1994 for an attempt to do so).

Here we pursue the previously stated goal of parsimony in theory development in science. In so doing, we examine the potential explanatory adequacy of the inhibitory reset hypothesis (principle 6) with respect to RB (see Bavelier & Jordan, 1992; Luo & Caramazza, 1996; MacKay & Miller, 1994 for similar analyses). It should be immediately noted, however, that Kanwisher herself initially entertained a similar account of RB but rejected it on the basis on one particular result. She observed that when participants do not have to respond to the first occurrence of the repeated item (only report the last item of the sequence), then the inhibitory effects of repetition turns to facilitation (Kanwisher, 1987, Exp. 3). However, both Park and Kanwisher (1994) and Luo and Caramazza (1996) point out that there has been some trouble in replicating this particular result (Kanwisher & Potter, 1990; Luo & Caramazza, 1996). One cannot therefore continue to use it as evidence against this particular account of RB.

Is RB, as obtained using the standard RSVP procedure, the result of the same mechanism that generates inhibitory above-threshold orthographic priming? We have argued in the present work that the inhibitory effects of above-threshold orthographically related primes reflect the operation of a reset mechanism that inhibits all activated representations when the target replaces the prime. However, in a recent paper, Chialant and Caramazza (1997) have argued that the inhibitory effects of orthographically similar items observed in the standard RB paradigm do not have the same underlying cause as the effects obtained with identical repetition. These authors measured RB for identical repetition (e.g. barn-barn) and for orthographically similar words (e.g. yarn-barn) at three different temporal lags (131, 261, and 383 ms) corresponding to one–three intervening items. The results show stronger inhibition with identical repetition at the shortest lag, similar effects for both conditions at the intermediate lag, and stronger inhibition for the orthographically related pairs at the longest lag. Thus, the results of Chialant and Caramazza (1997) show that inhibitory effects of identical repetition are strongest at first, and effects across orthographic neighbours become stronger with longer delays. The authors interpret the different time course of recall performance in these two conditions as indicating different underlying causes. However, if the degree of inhibition generated by the reset mechanism is a function of the activation level of a given representation at the moment of reset (principle 6), then identical repetition should suffer stronger inhibition than orthographically related items. In order to capture the fact that orthographically related items eventually suffer more inhibition than identical repetition, one must assume an inhibitory "rebound" function (cf. Houghton & Tipper, 1994), the duration of which is inversely proportional to the amplitude (i.e. the stronger the inhibition, the faster it develops and then dissipates). In this case, identical repetition causes fast-acting but short-lived inhibition, whereas orthographically similar items cause slower-acting more sustained inhibition.

According to principle 6 of our model, the inhibitory reset is triggered by a mismatch criterion that evaluates the dissimilarity of two successive stimuli. One prediction of this account applied to the RB phenomenon is that if the item intervening between repetitions is a non-linguistic pattern mask then recall of repeated items should be enhanced relative to a control condition. In this condition, no reset will be triggered and facilitatory repetition priming will occur. To our knowledge, this condition has not been tested with word stimuli. Park and Kanwisher (1994) obtained RB for letter repetitions with non-alphabetic symbols as intervening items. However, symbols such as ? and @ could lead to the partial activation of similar letters, thus triggering the reset mechanism. Park and Kanwisher also tested a condition with no item intervening between repetitions (repeated items were separated by a blank of variable duration), and observed strong effects of RB. Following the reasoning of Humphreys et al. (1988), however, when no stimulus intervenes between briefly presented repeated items, then they will be integrated as a single perceptual event and only one item reported. Although this can itself be considered a form of RB, within the present theoretical framework it has a very different underlying cause than effects observed with intervening stimuli.

Further experimentation using the RSVP procedure is required to test this critical prediction of the reset hypothesis. As noted previously, in standard priming conditions the data are contradictory with respect to whether or not prime-target repetition with an intervening pattern mask produces facilitation or inhibition (Hochaus & Johnston, 1996; Humphreys et al., 1988). Furthermore, extending the repetition priming paradigm to include intervening (unrelated) words will allow a closer evaluation of the similarities and differences of repetition priming and RB phenomena.

THEORETICAL DEVELOPMENTS

In the preceding sections we have presented a tentative theoretical framework for analysing effects of orthographic priming in visual word recognition and related phenomena. This theoretical framework was expressed as a set of verbalized processing principles that characterize the basic operations involved. These principles were generated within the theoretical perspective of an interactive activation model of human information processing in general, and visual word recognition in particular (McClelland & Rumelhart, 1981). Although the verbal statement of these principles has proved useful in our analysis of the different phenomena observed in an orthographic priming situation, only quantitative tests of a fully implemented version will provide a definitive test of this approach. Principle 6 (the reset mechanism) of our theoretical framework remains to be implemented in the algorithmic model in order to provide quantitative predictions relative to orthographic priming with identified primes.

Nevertheless, the part of the model that has been implemented so far (Grainger & Jacobs, 1996; Jacobs & Grainger, 1992), has proved successful in accounting for orthographic priming effects obtained with briefly presented masked primes (see Grainger, 1992 for a discussion of our earlier work). Indeed, our discovery that an extended version (Jacobs & Grainger, 1992) of the original interactive activation model (McClelland & Rumelhart, 1981), could provide a very accurate account of masked orthographic priming effects (Segui & Grainger, 1990), was one of the major motivations for pursuing this line of theorizing. This discovery was important in that no other model of visual word recognition can provide a principled account of the same data.

Our discussion of temporal integration of information in visual word recognition was deliberately limited to an analysis of orthographic codes. Deliberately isolating the role of one specific type of information allows a more precise, systematic investigation of the elementary operations involved in word perception, upon which extended analyses (including phonology, morphology, and semantics) can be built. We are well aware of the fact that writing systems were developed as a convenient means of coding oral language (for storage and transmission, for example). In alphabetic systems, the individual elements of the alphabet correspond (often in a one-to-many, and many-to-one relation) to the elementary sounds of the spoken language. This means that variations in the orthographic (alphabetic) description of words is generally highly correlated with variations in their phonological description (e.g. as a string of phonemes). The degree of correlation is a function of the consistency of the spelling-to-sound and the sound-to-spelling mappings (Ziegler et al., 1996).

There is a large body of evidence today, clearly indicating that phonological codes have an early influence on the process of visual word recognition (e.g. Ferrand & Grainger, 1992; Lukatela, Lukatela, & Turvey, 1993; Perfetti & Bell, 1991; Van Orden, Pennington, & Stone, 1990; Ziegler & Jacobs, 1995). The theoretical framework described in the present work can be easily extended to incorporate this body of evidence. We have, until now, ignored the influence of phonological information in visual word recognition, as a deliberate research strategy that, on the one hand seeks a compromise between the principle of Occam's razor (do not multiply entities beyond necessity) and the need for increasingly complex algorithmic models of human cognition, and, on the other hand follows the strategems of nested and canonical modeling (see Grainger & Jacobs, 1996; Jacobs & Grainger, 1994). We have recently implemented an extended interactive activation model with phonological representations (Jacobs, Rey, Ziegler, & Grainer, 1998) as a formal, heuristic tool for developing and testing hypotheses concerning the role of phonological codes in the mapping of orthography to semantics during visual word recognition. Following the strategy of nested modelling, this work builds on our existing knowledge of orthographic processing. On-going and future research will help elucidate the complex interactions between orthographic, phonological, morphological, semantic, and syntactic processing that arise during printed word perception.

REFERENCES

Altarriba, J., & Soltano, E.G. (1996). Repetition blindness and bilingual memory: Token individuation for translation equivalents. *Memory and Cognition*, *24*, 700–711.

Antos, S.J. (1979). Processing facilitation in a lexical decision task. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 527–545.

Arguin, M., & Bub, D. (1995). Priming and response selection processes in letter classification and identification tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1199–1219.

Bavelier, D. (1994). Repetition blindness between visually different items: The case of pictures and words. *Cognition*, *51*, 199–236.

Bavelier, D., & Jordan, M.I. (1992). A dynamical model of priming and repetition blindness. In C.L. Giles, S.J. Hanson, & J.D. Cowan (Eds.), *Advances in neural information processing systems* (Vol. 5, pp. 879–886). San Mateo, CA: Morgan Kaufmann.

Bavelier, D., & Potter, M.C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 134–147.

Bavelier, D., Prasada, S., & Segui, J. (1994). Repetition blindness between words: Nature of the orthographic and phonological representations involved. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1437–1455.

Boucart, M. (1996). *La reconnaissance des objets*. Grenoble, France: Presses Universitaires de Grenoble.

Burton, A.M. (1994). Learning new faces in an interactive activation and competition model. *Visual Cognition*, *1*, 313–348.

Chialant, D., & Caramazza, A. (1997). Identity and similarity factors in repetition blindness: Implications for lexical processing. *Cognition, 63,* 79–119.

Colombo, L. (1986). Activation and inhibition with orthographically similar words. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 226–234.

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

Davis, C., & Forster, K.I. (1994). Masked orthographic priming: The effect of prime-target legibility. *Quarterly Journal of Experimental Psychology*, *47*, 673–698.

Dean, M.P., & Young, A.W. (1996). An item-specific locus of repetition priming. *Quarterly Journal of Experimental Psychology*, *49A*, 269–294.

De Groot, A.M.B. (1984). Primed lexical decision: Combined effects of the proportion of related prime-target pairs and the stimulus-onset asynchrony of prime and target. *Quarterly Journal of Experimental Psychology*, *36A*, 253–280.

Dell, G.S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, *93*, 283–321.

Dell, G.S., & O'Seaghdha, P.G. (1994). Inhibition in interactive activation models of linguistic selection and sequencing. In T. Carr & D. Dagenbach (Eds.), *Inhibitory processes in attention, memory, and language.* (pp. 409–453 New York: Academic press.

Evett, L.J., & Humphreys, G.W. (1981). The use of abstract graphemic information in lexical access. *Quarterly Journal of Experimental Psychology*, *33A*, 325–350.

Ferrand, L. (1996). The masked repetition priming effect dissipates when increasing the inter-stimulus interval: Evidence from word naming. *Acta Psychologica*, *91*, 15–25.

Ferrand, L., & Grainger, J. (1992). Phonology and orthography in visual word recognition: Evidence from masked nonword priming. *Quarterly Journal of Experimental Psychology*, *42A*, 353–372.

Ferrand, L., & Grainger, J. (1993). The time-course of orthographic and phonological code activation in the early phases of visual word recognition. *Bulletin of the Psychonomic Society*, *31*, 119–122.

Ferrand, L., & Grainger, J. (1994). Effects of orthography are independent of phonology in masked form priming. *Quarterly Journal of Experimental Psychology*, *47A*, 365–382.

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*, 309–346.

Forster, K.I. (1976). Accessing the mental lexicon. In R.J. Wales & E.W. Walker (Eds.), *New approaches to language mechanisms* (pp. 257–287). Amsterdam: North-Holland.

Forster, K.I. (1987). Form priming with masked primes: The best match hypothesis. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 127–146). Hove, UK: Lawrence Erlbaum Associates Ltd.

Forster, K.I. (1993). Form-priming and temporal integration in word recognition. In G.T.M. Altmann & R. Shillcock (Eds.), *Cognitive models of speech processing: The second Sperlonga meeting* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Forster, K.I., & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 680–690.

Forster, K.I., & Davis, C. (1991). The density constraint on form-priming in the naming task: Interference effects from a masked prime. *Journal of Memory and Language*, *30*, 1–25.

Forster, K.I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *Quarterly Journal of Experimental Psychology*, *39A*, 211–251.

Giraudo, H., & Grainger, J. (in preparation). *Form and meaning in visual word recognition: An incremental priming study.*

Goldinger, S.D., Luce, P.A., & Pisoni, D.B. (1989). Priming lexical neighbours of spoken words: Effects of competition and inhibition. *Journal of Memory and Language*, *28*, 501–518.

Grainger, J. (1990). Word frequency and neighbourhood frequency effects in lexical decision and naming. *Journal of Memory and Language*, *29*, 228–244.

Grainger, J. (1992). Orthographic neighbourhoods and visual word recognition. In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology and meaning* (pp. 131–146). Amsterdam: North-Holland.

Grainger, J., & Ferrand, L. (1996). Masked orthographic and phonological priming in visual word recognition and naming: Cross-task comparisons. *Journal of Memory and Language*, *35*, 623–647.

Grainger, J., & Jacobs, A.M. (1993). Masked partial-word priming in visual word recognition: Effects of positional letter frequency. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 957–971.

Grainger, J., & Jacobs, A.M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, *103*, 518–565.

Grainger, J., & Jacobs, A.M. (1998). On localist connectionism and psychological science. In J. Grainger & A.M. Jacobs (Eds.), *Localist connectionist approaches to human cognition.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Greenwald, A.G., Klinger, M.R., & Schuh, E.S. (1995). Activation by marginally perceptible ("subliminal") stimuli: Dissociation of unconscious from conscious cognition. *Journal of Experimental Psychology: General*, *124*, 22–42.

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, *11*, 23–63.

Hamburger, M., & Slowiaczek, L.M. (1996). Phonological priming reflects lexical competition. *Psychonomic Bulletin and Review*, *3*, 520–525.

Hinton, J., Liversedge, S., & Underwood, G. (1996, September). *A competitive role for neighbours demonstrated using a partial priming methodology*. Paper presented at the 9th conference of the European Society for Cognitive Psychology, Würzburg, Germany.

Hochaus, L., & Johnston, J.C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 355–366.

Hochaus, L., & Marohn, K.M. (1991). Repetition blindness depends on perceptual capture and token individuation failure. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 422–432.

Houghton, G., & Tipper, S.P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language.* San Diego, CA: Academic Press.

Humphreys, G.W., Besner, D., & Quinlan, P.T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General*, *117*, 51–67.

Humphreys, G.W., Evett, L.J., & Quinlan, P.T. (1990). Orthographic processing in visual word identification. *Cognitive Psychology*, *22*, 517–560.

Humphreys, G.W., Evett, L.J., & Taylor, D.E. (1982). Automatic phonological priming in visual word recognition. *Memory and Cognition*, *10*, 128–152.

Jacobs, A.M., & Grainger, J. (1991). Automatic letter priming in an alphabetic decision task. *Perception and Psychophysics*, *49*, 43–52.

Jacobs, A.M., & Grainger, J. (1992). Testing a semistochastic variant of the interactive activation model in different word recognition experiments. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1174–1188.

Jacobs, A.M., & Grainger, J. (1994). Models of visual word recognition: Sampling the state of the art. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1311–1334.

Jacobs, A.M., Grainger, J. & Ferrand, L. (1995). The incremental priming technique: A method for determining within-condition priming effects. *Perception and Psychophysics*, *57*, 1101–1110.

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). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Kanwisher, N.G. (1987). Repetition blindness: Type identification without token individuation. *Cognition*, *27*, 117–143.

Kanwisher, N.G. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 404–421.

Kanwisher, N.G., & Potter, M.C. (1990). Repetition blindness: Levels of processing, *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 30–47.

Lukatela, G., Lukatela, K., & Turvey, M.T. (1993). Further evidence for the phonological constraints on visual lexical access: TOWED primes FROG. *Perception and Psychophysics*, *53*, 461–466.

Luo, C.R., & Caramazza, A. (1996). Temporal and spatial repetition blindness: Effects of presentation mode and repetition lag on the perception of repeated items. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 95–113.

Lupker, S.J., & Colombo, L. (1994). Inhibitory effects in form priming: Evaluating a phonological competition explanation. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 437–451.

MacKay, D.G. (1987). *The organization of perception and action: A theory of language and other cognitive skills.* New York: Springer.

MacKay, D.G., Abrams, L., Pedroza, M.J., & Miller, M.D. (1996). Cross-language facilitation, semantic blindness, and the relation between language and memory: A reply to Altarriba and Soltano. *Memory and Cognition*, *24*, 712–718.

MacKay, D.G., & Miller, M.D. (1994). Semantic blindness: Repeated concepts are difficult to encode and recall under time pressure. *Psychological Science*, *5*, 52–55.

Marcel, A.J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, *15*, 197–237.

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

Massaro, D.W., & Cohen, M.M. (1991). Integration versus interactive activation: The joint influences of stimulus and context on perception. *Cognitive Psychology*, *23*, 558–614.

McClelland, J.L., & Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception: P. 1. An account of basic findings. *Psychological Review*, *88*, 375–407.

Meiran, N. (1996). Is reading ability related to activation dumping speed? Evidence from immediate repetition priming. *Memory and Cognition*, *24*, 41–59.

Monsell, S. (1985). Repetition and the lexicon. In A.W. Ellis (Ed.), *Progress in the psychology of language, Vol. 2.* Hove, UK: Lawrence Erlbaum Associates Ltd.

Monsell, S. (1991). The nature and locus of word frequency effects in reading. In D. Besner & G.W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 148–197). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, *76*, 165–178.

Naish, P. (1980). The effects of graphemic and phonemic similarity between targets and masks in a backward visual masking paradigm. *Quarterly Journal of Experimental Psychology*, *32*, 57–68.

Napps, S. (1989). Morphemic relationships in the lexicon: Are they distinct from semantic and formal relationships? *Memory and Cognition*, *17*, 729–739.

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

Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *106*, 226–254.

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G.W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Park, J., & Kanwisher, N. (1994). Determinants of repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 500–519.

Peressotti, F., & Grainger, J. (1995). Letter-position coding in random consonant arrays. *Perception & Psychophysics*, *57*, 875–890.

Peressotti, F., & Grainger, J. (in press). The role of letter identity and letter position in orthographic priming. *Perception & Psychophysics.*

Perfetti, C.A., & Bell, L. (1991). Phonemic activation during the first 40 ms of word identification: Evidence from backward masking and priming. *Journal of Memory and Language*, *30*, 473–485.

Perfetti, C.A., Bell, L.C., & Delaney, S.M. (1988). Automatic (prelexical) phonetic activation in silent word reading: Evidence from backward masking. *Journal of Memory and Language*, *27*, 59–70.

Ratcliff, R., Hockley, W., & McKoon, G. (1985). Components of activation: Repetition and priming effects in lexical decision and recognition. *Journal of Experimental Psychology: General*, *114*, 435–450.

Rueckl, J.G. (1990). Similarity effects in word and pseudoword repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 374–391.

Rumelhart, D.E., & McClelland, J.L. (1982). An interactive activation model of context effects in letter perception: P. II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, *89*, 60–94.

Rumelhart, D.E., & Norman, D.A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science*, *6*, 1–36.

Segui, J., & Grainger, J. (1990). Priming word recognition with orthographic neighbours: Effects of relative prime-target frequency. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 65–76.

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

Slowiaczek, L.M., & Hamburger, M. (1992). Prelexical facilitation and lexical interference in auditory word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1239–1250.

Stone, G.O., & Van Orden, G.C. (1993). Strategic control of processing in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 744–774.

Stone, G.O., & Van Orden, G.C. (1994). Building a resonance framework for word recognition using design principles and system principles. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1248–1268.

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

Van Orden, G.C., Pennington, B.F., & Stone, G.O. (1990). Word identification in reading and the promise of subsymbolic psycholinguistics. *Psychological Review*, *97*, 488–522.

Ziegler, J.C., & Jacobs, A.M. (1995). Phonological information provides early sources of constraint in the processing of letter strings. *Journal of Memory and Language*, 34, 567–593.

Ziegler, J.C., Jacobs, A. M., & Stone, G.O. (1996). Statistical analysis of the bi-directional inconsistency of spelling and sound in French. *Behavior Research Methods, Instruments, and Computers*, *28*, 504–515.