



ACADEMIC
PRESS

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Memory and Language 48 (2003) 207–232

Journal of
Memory and
Language

www.elsevier.com/locate/jml

Strategic control over rate of processing in word reading: A computational investigation

Christopher T. Kello^{a,*} and David C. Plaut^b

^a Department of Psychology, George Mason University, Fairfax, VA, USA

^b Departments of Psychology and Computer Science, Center for the Neural Basis of Cognition, Carnegie Mellon University, USA

Received 11 December 2001; revision received 15 April 2002

Abstract

It has recently been proposed that the time course of lexical processing can be influenced by strategic control factors in word reading. In the current study, a specific hypothesis of control, rate of processing, was implemented in three simulations of word reading. Simulation results showed that, by modulating the dynamic of processing, the control parameter *input gain* can account for effects of pressure for speed and stimulus blocking on naming performance. Results showed also that, to account for error patterns found in the tempo-naming task (Kello & Plaut, 2000), the influence of lexical knowledge must be strengthened relative to the standard dual-pathway architecture. Two methods of strengthening the influence of lexical knowledge are demonstrated: input gain used as a mechanism of route emphasis, or integration of the two commonly proposed pathways from orthography to phonology.

© 2002 Elsevier Science (USA). All rights reserved.

Keywords: Word reading; Strategic control; Processing rate; Input gain; Connectionist model; Timing

What we know about word reading is largely based upon empirical studies that examine how reading behavior, both normal and impaired, is influenced by characteristics of printed stimuli. These stimulus factors have been, for the most part, physical (e.g., stimulus discriminability and confusability), lexical and sub-lexical (e.g., printed frequency and spelling–sound consistency), or relational (e.g., semantic or phonological similarity between a prime and target). However, knowledge about word reading has also been acquired by examining how reading behavior is influenced by factors beyond the stimuli themselves. Such *extra-stimulus* factors include explicit control mechanisms such as those that might be invoked by the instructions given to a participant, as well as implicit control mechanisms such as those that might respond to the overall distribution of stimuli in a given experiment (e.g., Baluch &

Besner, 1991; Colombo & Tabossi, 1992; Jared, 1997; Kello & Plaut, 2000; Lupker, Brown, & Colombo, 1997; Monsell, Patterson, Graham, Hughes, & Milroy, 1992; Paap & Noel, 1991; Rastle & Coltheart, 1999; Stone & VanOrden, 1993).

Given that stimulus factors have played the primary role in empirical research on word reading, it is not surprising that theories and computational models of word reading are most well-developed with regard to these factors. By comparison, researchers have only just begun to formulate and simulate mechanisms that may account for the role of extra-stimulus factors in word reading (for an example, see Zorzi, 1999). Moreover, any control mechanism proposed to account for findings in word reading must eventually be integrated within a general theory of the cognitive processes involved in reading.

In the current study, we investigated a general mechanism of strategic control, termed *input gain* (Cohen, Dunbar, & McClelland, 1990; Kello & Plaut, 2000; Kello, Plaut, & MacWhinney, 2000), in the context of word reading. Input gain is a control parameter that globally

* Corresponding author. Fax: +703-993-1330.

E-mail address: ckello@gmu.edu (C.T. Kello).

scales the magnitudes of net inputs to units within a connectionist network. One implication of changes in gain is to alter the *rate of processing* within the network. Our primary goal in the current work was to test the hypothesis that manipulations of input gain can account for changes in naming behavior that have been observed empirically as a result of extra-stimulus influences on the time course of responding. In light of this purpose, we evaluated our computational results with two broad goals in mind: (1) demonstrate how input gain can modulate processing in connectionist models of word reading and (2) test input gain as a general mechanism of control over the cognitive processes involved in word reading.

The naming behavior that served as an empirical benchmark for our computational investigations came from two experimental paradigms. First, Kello and Plaut (2000) examined control over response initiation in word reading by instructing subjects to initiate naming in time with a rhythmic cue—a task that was termed *tempo-naming*. Tempo-naming provides an explicit, extra-stimulus cue to manipulate the time course of responding. Second, Jared (1997) and Lupker et al. (1997) reported evidence that the proportion of relatively difficult or easy stimuli in a list or block can affect naming latencies to those stimuli (further evidence was reported in Taylor & Lupker, 2001). In this case, evidence suggests that stimulus blocking can serve as an implicit, extra-stimulus influence on the time course of responding.

We examined the effects of input gain in three connectionist simulations of word reading, and compared them against a set of empirical results from the tempo-naming and stimulus blocking paradigms. All three simulations showed that, under manipulations of input gain, connectionist models of word reading can simulate a number of specific patterns of behavior found in the tempo-naming and stimulus blocking paradigms. These computational results supported the hypothesis that input gain provides a basis for strategic control in word reading, by altering rate of processing and the overall dynamics of processes within the reading system.

Strategic control in word reading

Extra-stimulus factors that affect cognitive processes have often been referred to as “strategic control” factors. In research on word reading, two hypothesized mechanisms of control have received the most attention: the *route emphasis* mechanism and the *time criterion* mechanism. Here we explain the workings of each mechanism and briefly review the supporting evidence. This discussion serves as context for a third candidate mechanism of control, rate of processing.

Route emphasis is currently the most widely accepted hypothesis of control in word reading (Baluch & Besner, 1991; Colombo & Tabossi, 1992; Coltheart & Rastle,

1994; Herdman, 1992; Monsell et al., 1992; Tabossi & Laghi, 1992; Zevin & Balota, 2000). This hypothesis is a natural extension of two competing theoretical frameworks in word reading: the dual-route and “triangle” frameworks. In both theoretical frameworks, the ability to name a printed word or pseudoword (i.e., a pronounceable string of letters that does not spell a word) is supported by at least two processing routes. In the dual-route framework (e.g. Coltheart, Curtis, Atkins, & Haller, 1993; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001), a *sub-lexical* route stores grapheme–phoneme correspondence rules, whereas a *lexical* route stores information connecting the printed forms of known words to their respective pronunciations. In the triangle framework (e.g., Harm & Seidenberg, 1999; Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989), word reading processes take the form of cooperative and competitive interactions among orthographic, phonological, and semantic representations (via learned “hidden” representations that mediate among each pair of these). Thus, this framework also has two pathways from print to sound: a *phonological* pathway that maps orthography directly to phonology, and a *semantic* pathway that maps orthography to phonology via semantics. Although the nature of computation within the dual-route and triangle frameworks is very different, they share the assumption that one route or pathway captures sub-lexical spelling–sound correspondences whereas the other captures lexical/morphological correspondences.¹ Without intending to deny the importance of differences between the two underlying theoretical perspectives, for clarity, we will refer to this distinction using the terms “sub-lexical” and “lexical” routes, respectively.

Given the separate and complementary nature of these routes, it is natural to wonder whether the contribution of one route to processing might be increased or decreased depending on task demands. With respect to the task of word naming, route emphasis is the hypothesis that the contribution of one route to the computation of phonology can be emphasized or de-emphasized relative to the contribution of the other route. This hypothesis is motivated by the idea that certain reading tasks would seem to benefit from the ability to selectively emphasize one type of processing or the other. For example, the lexical route is better suited to processing words—and, indeed, is critical for correctly pronouncing *irregular*

¹ In fact, Zorzi, Houghton, and Butterworth (1998) have developed a “dual-process” connectionist model that has strong commonalities with both the dual-route and triangle frameworks. The sublexical pathway in the model captures grapheme–phoneme correspondences using direct connections between orthography and phonology (without hidden units), whereas the lexical pathway (like that in the dual-route framework) uses localist word units.

words like PINT that violate standard spelling–sound correspondences (SSCs)—and can sometimes interfere with the processing of nonwords. The sub-lexical route is better suited to processing nonwords, and can sometimes interfere with word processing.

Although the intrinsic motivation may be compelling, one would also want empirical support for route emphasis before adopting it as a working hypothesis. Researchers have, in fact, tested the hypothesis, and results have generally been consistent with a route emphasis account. The main source of evidence for the route emphasis hypothesis has come from blocking studies that manipulate the proportion of stimuli hypothesized to interfere with one or the other processing route (Andrews, 1982; Frederiksen & Kroll, 1976; Monsell et al., 1992; Rastle & Coltheart, 1999). The logic behind these studies is as follows: if most or all stimuli in a given block (i.e., a *pure* block) are better suited to one type of processing, then readers should emphasize that type of processing (and de-emphasize any interfering processes), if possible. The effect of route emphasis should be observable if performance in the pure block is compared with performance in a *mixed* block containing various types of stimuli.

To give an example, Monsell et al. (1992) divided stimuli in a word naming task into pure and mixed blocks. The pure blocks contained either all nonwords or all irregular words (of mixed frequency in Experiment 1, and separated by frequency in Experiment 2). The mixed blocks contained both nonwords and irregular words. Monsell and his colleagues found that irregular words were generally named faster in pure versus mixed blocks. The authors interpreted their results as evidence that subjects de-emphasized the sub-lexical route in pure blocks of irregular words to reduce interference from sub-lexical processing.

Although some results from stimulus blocking studies have been interpreted as support for the route emphasis hypothesis, certain findings have proven difficult to reconcile with a route emphasis account. Most notably, Monsell et al. (1992) found that the latency advantage for irregular words in pure blocks (compared to blocks mixed with nonwords) was reliable only for pure blocks of high-frequency (HF), but not low-frequency (LF), words (for similar results, see Andrews, 1982; Frederiksen & Kroll, 1976). Lupker et al. (1997) and Jared (1997) revisited the stimulus blocking results, and noted that if one defines “de-emphasis” as slowed processing times (of the sub-lexical or phonological route in this case), then LF irregular words should have an equal or greater advantage in the pure block compared with HF irregular words. This is because processing times for nonwords must overlap more with LF compared to HF words, provided that the mean processing time of the sub-lexical route is greater than that of the lexical route (as suggested by findings such as the lexicality advan-

tage; Forster & Chambers, 1973). By contrast, studies have found a greater pure block advantage for HF irregular words.

Lupker et al. (1997) reran the Monsell et al. (1992) blocking experiment (with minor variations) and replicated the pure block advantage for HF irregular words. Moreover, they found a statistically reliable pure block *disadvantage* for LF irregular words, whereas the route emphasis hypothesis predicts a pure block advantage. Lupker et al. (1997) ran a second experiment in which all of the stimuli contained regular mappings from spelling to sound. In this case, no blocking effect was expected on the basis of route emphasis because the sub-lexical route should remain active in both the pure and mixed blocks. Contrary to expectations, a pure block advantage was found for HF words (now regular) as well as for LF words. Jared (1997) found analogous results, except that she compared blocks mixed with nonwords versus blocks mixed with LF inconsistent words. In summary, the results from Jared (1997) and Experiments 1 and 2 from Lupker et al. (1997) were inconsistent with the route emphasis hypothesis because, according to route emphasis, mixing with nonwords should not speed LF irregular words nor influence regular words.

To explain their results, (Lupker et al., 1997) re-categorized stimuli as fast or slow on the basis of mean latencies in the pure blocks. The nonwords and LF irregular words were categorized as slow, and the HF words were categorized as fast (LF regular words were in the middle). By labeling stimuli in this way, the pattern of results could be described as follows: whenever fast and slow stimuli were mixed, response latencies increased for the fast stimuli, but decreased for the slow stimuli, relative to when those stimuli were placed in pure blocks. This insight led Lupker and his colleagues to propose that the blocking manipulation prompted subjects to adjust a time criterion to initiate naming responses. The general idea was that subjects set a time deadline relative to stimulus onset (Ollman & Billington, 1972). If a pronunciation is not fully activated by the deadline, then the response is initiated on the basis of whatever information is available at that time (also see Meyer, Osman, Irwin, & Kounios, 1988). To maintain a certain level of accuracy while responding quickly, subjects adjust the time criterion based on the difficulty of the stimuli presented during the experiment. A pure block of fast stimuli allows for an earlier time criterion compared with a pure block of slow stimuli. When fast and slow stimuli are mixed, subjects set a middling time criterion: thus, fewer HF but more LF words are hurried.

Control in the tempo-naming task

The time criterion is a mechanism that, in principle, may enable fairly tight control over the amount of time

available for processing a printed stimulus. Of course, tight control does not necessarily mean total control: an overly simple time criterion would suggest that responses can be initiated precisely at the time criterion setting, regardless of the status of stimulus processing. Such an absolute mechanism would seem unlikely given the existence of stimulus effects, but nonetheless, the mechanism raises an interesting question: to what extent can readers be explicitly instructed to control the amount of time available for processing?

Kello and Plaut (2000) introduced the tempo-naming task partly as a means of addressing this question. The main idea behind the tempo-naming task is to take advantage of people's natural ability to entrain behavior to an external rhythm. The task is structured as follows. Each trial begins with a sequence of auditory beeps that are timed at a specific tempo. Each beep (i.e., beat of the tempo) is accompanied by the removal of a visual flanker on the computer screen (see Fig. 1). A target stimulus is presented upon the final beep, and the participant is instructed to pronounce the target such that the response is initiated simultaneously with the subsequent beep (which is not actually presented). After each response, feedback is given on how well the response was timed with the rhythm.

The tempo provides an explicit and precise cue to control the amount of time from stimulus onset to response initiation. To the extent that the cue is effective, faster tempos should shorten the time course whereas slower tempos should lengthen it. What is unclear in this task is what kind of cognitive mechanism(s) might be used to accomplish it, but a time criterion is one possibility. If a strict time criterion could be set to the rhythm, responses would essentially be timed with

the rhythm (modulo articulatory and acoustic variability).

The time criterion account of performance in the tempo-naming task raises two issues. First, to the extent that responses can be well-timed with the rhythm, stimulus effects on naming latencies should be attenuated. Second, to the extent that responses can be well-timed with very fast tempos (i.e., less than 500 ms), those fast responses should reflect the premature results of processing. By "premature," we mean that the time criterion would need to be set very early in the time course of processing, and a response would therefore need to be initiated on the basis of whatever processing is complete at that time.

Kello and Plaut (2000) explored these issues by comparing naming performance in the standard naming task with that of the tempo-naming task. The results of three experiments with the tempo-naming task can be summarized as follows. Participants were largely able to entrain response initiation to an external tempo, and the fastest tempo interval (150 ms faster than each subject's mean latency in a standard naming task) drove response latencies to be about 100 ms faster than baseline, on average. Faster tempos also induced a speed/accuracy trade-off. The effects of printed frequency and spelling–sound consistency were attenuated in the tempo-naming task compared with those found in a standard naming task. These results were consistent with the time criterion hypothesis, but two additional results were not.

First, faster tempos caused response durations (as measured by time from acoustic onset to acoustic offset) to decrease. Second, as tempo increased, the rate of spelling–sound errors remained constant whereas the rates for other types of errors increased. Spelling–sound

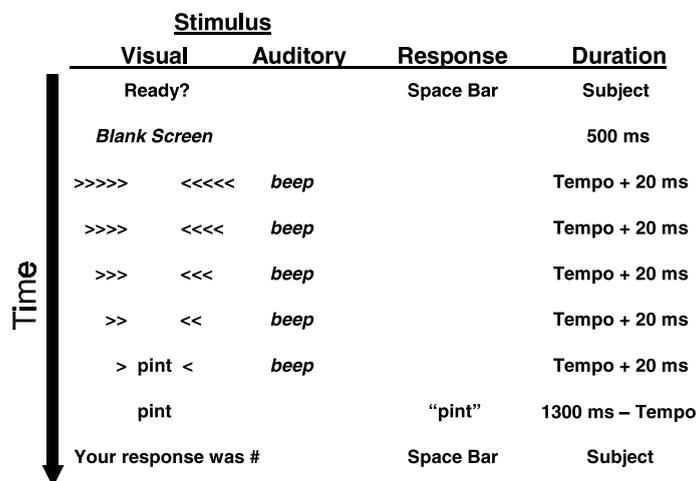


Fig. 1. Diagram of the course of events for a single trial in the tempo-naming task. The "> <" symbols are flankers indicating the position of the target stimulus. Tempo is the time interval between each beep, determined by the tempo condition and the subject's baseline. "Subject" indicates that the duration is subject-dependent.

errors were defined as legitimate alternative pronunciations of a letter string based on its constituent orthographic units (i.e., Legitimate Alternative Reading of Components, or LARC errors; Strain, Patterson, Graham, & Hodges, 1998).² LARC errors that formed a word, such as *SEW* pronounced as *SUE*, were classified as *mixed* errors. Other error types were (non-LARC) word errors, nonword errors, and articulatory errors. Nonword pronunciations were orthographically/phonologically similar to their targets, but did not form words. Articulatory errors included stutters, mis-starts, and failures to respond.

By itself, the time criterion hypothesis does not make predictions concerning response durations or error patterns. A time criterion (in combination with one or more activation criteria) governs response initiation only, not response execution. Response durations and errors are the results of response execution, not response initiation. Response execution is a product of the interaction between response criteria and processes in the word reading system. Therefore, to address naming durations and errors, one would need to consider the word reading system as a whole. Kello and Plaut (2000) made an initial foray into this issue by examining the effect of a time criterion in two competing models of word reading: the dual-route cascaded (DRC) model (Coltheart et al., 1993) and the distributed attractor model reported in Plaut et al. (1996). Both simulations failed to account for the result that, as tempo increased, the number of LARC errors remained constant whereas the occurrence of other types of errors increased. The simulations produced a mostly equal increase in all types errors as the time criterion was shifted to successively earlier points in processing.

The rate of processing hypothesis

The absence of an account for two aspects of the tempo-naming results prompted us to consider an alternative hypothesis. The observed effect of tempo on response durations and error patterns suggested that the tempo cues did not simply influence a response criterion but affected processing within the word reading system itself. In particular, the effect of tempo on response durations suggested to us that the *rate of processing* was affected by the tempo cues. Kello et al. (2000) made a connection between response durations and rate of processing on the basis of *cascaded articulation*. Cascaded articulation is the hypothesis that, under some circumstances, response processing and response execution can overlap in time (also see Ka-

wamoto, Kello, Jones, & Bame, 1998; Kawamoto, Kello, Higareda, & Vu, 1999). If articulation was cascaded in the tempo-naming study, then one would expect a change in the rate of response processing to be reflected in a change in the rate of response execution. If one treats response duration as a coarse measure of the rate of response execution, then the observed decrease in response duration with faster tempos indicated an increase in the rate of processing. However, this logic requires that articulation was, in fact, cascaded in the tempo-naming task. Kello et al. (2000) showed that articulation can become cascaded with processing under conditions in which speeded responding is emphasized. Given that speeded responding was emphasized in the tempo-naming task, it seems reasonable to assume that articulation was cascaded in the tempo-naming study.

The rate of processing hypothesis raises at least two questions: what was the operative mechanism of control over rate of processing in the tempo-naming experiments, and how might this mechanism account for the tempo-naming results? We address these questions in the next two sections.

Input gain as a mechanism of control over rate of processing

The rate of processing hypothesis, as stated thus far, requires that a specific mechanism of control is proposed to implement the hypothesis in a computational model. In the current study, we investigated *input gain* as a specific mechanism of control to implement the rate of processing hypothesis. Input gain is a multiplicative scaling parameter on the net inputs to connectionist processing units (equivalent to the inverse of *temperature* in Boltzmann machines; Ackley, Hinton, & Sejnowski, 1985). Input gain can be thought of as controlling a unit's sensitivity to inputs from other units. Under low levels of input gain, net inputs must be large in magnitude to produce outputs that are large in magnitude. Under high levels of input gain, smaller magnitudes of net input are sufficient to produce large outputs. In this section, the effect of input gain is demonstrated in five simple cases. The purpose of these demonstrations is two-fold: to illustrate how input gain modulates processing in connectionist processing units, and to show how, in principle, input gain might account for some of the basic results from the tempo-naming task.

In Fig. 2, the output of a standard connectionist processing unit is plotted as a function of net input and time, for different values of input gain. The activation function used for this illustration was the logistic,

$$a_j^{[l]} = (1 + \exp(-x_j^{[l]}\gamma))^{-1}, \quad (1)$$

² For example, pronouncing *PINT* to rhyme with *MINT*, *MOU* to rhyme with *NOW*, or *NOW* to rhyme with *MOU*.

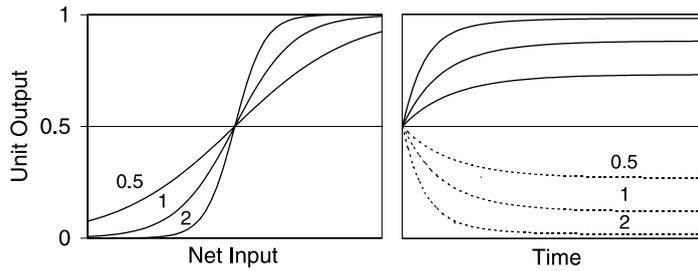


Fig. 2. In the left graph, unit activation is plotted as a function of net input and input gain at a fixed time t . In the right graph, unit activation is plotted as a function of time and input gain (0.5, 1, or 2) at a fixed net input of 2 (solid lines) or -2 (dashed lines).

where $a_j^{[t]}$ is the activation of unit j at time t , $x_j^{[t]}$ is the net input to unit j at time t , and γ is input gain. For purposes of computation, time was discretized into a series of ticks t with duration τ . The net input to a given unit j at time t was calculated as

$$x_j^{[t]} = \tau I_j^{[t]} + (1 - \tau)x_j^{[t-\tau]}, \quad (2)$$

where the net input $I_j^{[t]}$ was calculated as

$$I_j^{[t]} = \sum_i w_{ij} a_i^{[t-\tau]} + E_j^{[t]}, \quad (3)$$

where w_{ij} was the weight on the connection from unit i to unit j and E_j was the external input to unit j .

To give unit processing a temporal extent, either unit inputs or unit outputs can be calculated as a weighted average of the current state and previous states (i.e., *time averaging*). In the current example, as well as for all unit computations in the current study, unit inputs were time averaged (for a discussion of the difference between time averaged inputs and outputs, as well as theoretical motivation for the use of time averaged inputs, see Harm & Seidenberg, 1999).

Fig. 2 shows that as the net input moves away from zero, unit activation moves from 0.5 towards an asymptote at zero or one (for increasingly negative or positive net inputs, respectively). Input gain modulates the sharpness of this function such that higher values of gain cause the function to become more binary—the asymptote moves closer to zero or one and unit activation approaches zero or one more quickly as net input moves away from zero. An increase in input gain has a similar effect on unit output as a function of time (with the net input held constant).

This simple illustration shows that, for a single unit, input gain provides a mechanism of control over rate of processing. However, models of cognitive processes require networks of units. We must therefore extend our analysis of input gain to a network. The influence of input gain on any given unit in a network is potentially more complex than the single unit case. This is because input gain can affect not only the scaling of net inputs,

but the net inputs themselves. The effect of input gain is determined by the pattern of network connectivity and external inputs, as well as the activation function and model of time (among other factors). A full analysis of all these factors is a topic of study in itself, and we do not attempt to provide such an analysis here. To illustrate the range of effects that input gain can have, a network of three units, i , j , and k , was run with four different patterns of connectivity. For all patterns of connectivity, external input was fixed at 2 on unit i and 0 on units j and k . Net inputs were initialized to zero and τ was fixed at 0.1. The results are shown in Fig. 3.

Input gain had four different effects on the time course of activation for unit k , depending the pattern of connectivity among the units.

- Similar to the single unit case, asymptote moved toward one more quickly and more closely with higher levels of input gain.
- Asymptote changed from one to 0.5 with higher levels of input gain.
- Asymptote changed from one to zero with higher levels of input gain.
- As the level of input gain increased, the output of unit k began to oscillate, and the oscillations subsided over time. Higher levels of input gain increased the amplitude and duration of oscillation.

Expected effects of input gain

The demonstrations of the previous section are not meant to represent all of the possible effects of input gain, nor do they imply that each effect is equally likely. For example, with static external input, recurrent connections are necessary to induce oscillatory behavior, so behavior (d) will not be observed in feed-forward networks. Also, this network has only three units, so one might expect a wider range of behaviors in more complex networks. These examples demonstrate that input gain can have variable effects on network behavior, depending on the network's configuration. This variability suggests that input gain, as a hypothesized mechanism

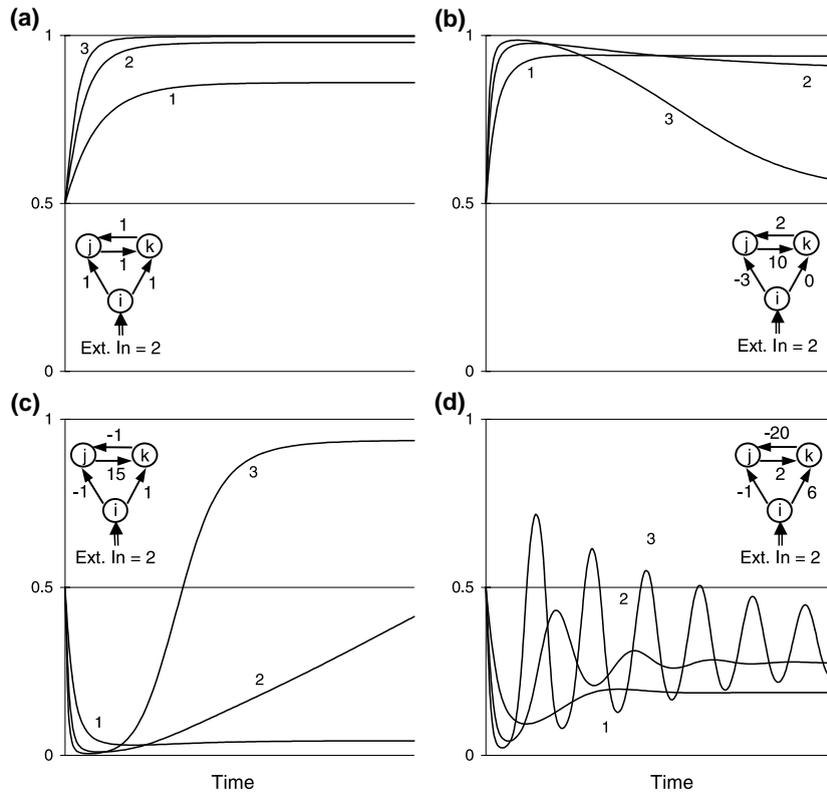


Fig. 3. The output for unit k is plotted as a function of time and input gain for four different weight configurations of the illustrated network. Unit output is plotted under three different levels of input gain (1, 2, and 3) for each network (trajectories are labeled). Weight configurations are shown in each panel.

of control over rate of processing in word reading, must be tested in specific models of word reading. The main purpose of the current study was to conduct such tests, but before these are reported, it would be useful to briefly consider how input gain might, in principle, account for the targeted empirical results. The targeted results come from the two experimental paradigms that were discussed in previous sections: tempo-naming and stimulus blocking.

With regard to the tempo-naming experiments (Kello & Plaut, 2000), we targeted four results: as tempo increased, (1) naming latencies and durations decreased, (2) error rates increased, (3) LARC and mixed errors showed little or no increase, and (4) of all error types, word errors were most frequent. One might expect the first result to be accounted for if the time course of response execution is in some way coupled to the time course of processing in a model of word reading. As demonstrated for a single connectionist processing unit, increased levels of input gain can cause unit outputs to reach asymptote more closely and more quickly. Thus, the time course of processing can be considered to become compressed under high levels of input gain. If the

time course of processing is coupled to the time course of response execution, then one would expect latencies and durations to become compressed as well. As explained in the Rate of Processing hypothesis, evidence for cascaded articulation gives us reason to assume that processing is, in fact, coupled with execution to some degree.

With regard to the second result, the demonstrated effects of input gain in a network provide some evidence that higher levels of input gain can cause an increase in response errors. In particular, in the network configurations shown in panels (b), (c), and (d) of Fig. 3, the asymptotic output of a unit diverged as a function of input gain. If target responses are considered to be in the corners of unit output space (as they typically are), then a change in asymptotic output would cause a change in response. If correct responses are produced under “normal” levels of input gain (i.e., those that the network is trained under), then a change in asymptotic response would cause an error.

With regard to the third and fourth results, we can interpret LARC and mixed errors as arising from an erroneous influence of SSCs, whereas word errors arise

from an erroneous influence of lexical knowledge. On the basis of this interpretation, higher levels of input gain would need to attenuate the influence of SSCs on response processing, relative to the influence of lexical knowledge (for other relevant results, see Farrar, Orden, & Hamouz, 2001; Hendriks & Kolk, 1997). The demonstrations in the current section do not shed light on how input gain might cause such a result because the example networks were not models of word reading. The simulations reported in the following sections address this issue.

Before the simulations are reported, two additional results from the tempo-naming experiments need to be addressed. First, the effect of stimulus factors (i.e., word frequency, spelling–sound consistency, and lexicality) was attenuated in the tempo-naming task relative to the standard naming task. The reason for this effect is presumably because participants were instructed to respond at a very specific point in time, regardless of the nature of the stimulus. On our approach, the precise nature of response timing in the tempo-naming task requires a timing mechanism in addition to a rate of processing mechanism. On each trial, rate of processing is adjusted in accordance with the rate of tempo, and a timing set point is aligned in accordance with the temporal locations of the tempo cues.

Therefore, to account for the attenuation of stimulus effects in the current models, a timing mechanism would need to be implemented. Input gain would set the rate of processing at level that is generally appropriate for a given tempo, and the timing mechanism might, for example, delay any responses that reach criterion early so that they are initiated in time with the tempo.

We chose not to implement a timing mechanism for four main reasons. First, given that the timing mechanism would have a very small influence on naming behavior compared with input gain, it is unlikely to have any impact on the nature of the responses generated by a model. Second, if both an input gain and a timing mechanism were implemented, it would be difficult to attribute the cause of simulation results to input gain versus the timing mechanism. Third, without perceptuo-motor constraints, the degrees of freedom in implementing a simple timing mechanism (i.e., its initial set point and decrement with each step of increase in input gain) would be too unconstrained. Finally, we hypothesized that input gain would, in fact, at least partially account for the attenuation of stimulus effects. The modulation of stimulus effects by input gain is reported in the Results of Simulation 1.

The second result that needs to be addressed is that, in the tempo-naming experiments, the onset of voicing seemed to be the articulatory/acoustic event that was timed with the tempo. This conclusion was based on the way that timing was influenced by the articulatory/acoustic characteristics of the initial phoneme of each

response. As with the attenuation of stimulus effects, an account of this result would require the implementation of a timing mechanism. We did not address this result because we did not implement a timing mechanism.

In addition to the tempo-naming experiments, we targeted results from the stimulus blocking experiments reported by Jared (1997) and Lupker and his colleagues (Lupker et al., 1997; Taylor & Lupker, 2001). In particular, for any given stimulus, latencies were found to be slower if that stimulus is mixed with more difficult stimuli (in terms of stimulus–response processing), and faster if the stimulus is mixed with relatively easier stimuli. The time criterion hypothesis explains this general effect by positing an internal response deadline that is adjusted to balance speed and accuracy. On the basis of the demonstrations in the current section, we propose a similar account using input gain as the mechanism of control, instead of a time criterion. If more difficult stimuli are more error prone under high levels of input gain, then blocks of difficult stimuli would require moderate levels of input gain to keep errors at a minimum. Conversely, blocks of relatively easy stimuli would encourage higher levels of input gain, provided that the task requires fast and accurate responses.

Modeling approach

In implementing the current simulations, there were a number of modeling decisions to make, some more relevant to the issues at hand than others. There are at least two general approaches to cognitive modeling that offer some guidance in making such decisions. We refer to one as the *realist* approach and the other as the *fundamentalist* approach.

The realist approach holds that a model should encompass as much scale and detail as possible of the phenomenon it proposes to explain. Proponents of this approach argue that “nature is smarter than we are”—that it is difficult to reason what aspects of a phenomenon are crucial for an explanation. Therefore, the model should encompass as much of the phenomenon as possible to avoid neglecting important details. It is also argued that if a model reduces the size and complexity of the phenomenon, one does not know if that model will successfully scale up to the true size of the phenomenon.

The fundamentalist approach holds that, as much as possible, a model should embody only the principles that are theorized to account for the phenomenon in focus, and extraneous details should be abstracted out of the model. Proponents of this approach argue that too much detail obscures the relationship of the model to the theory. Relatedly, a realist model can become an empirical phenomenon in its own right, subject to empirical investigation to determine why it did or did not account for a set of results. Pragmatically speaking, large and

complex models can require a prohibitive amount of time and computing resources.

We primarily followed a fundamental approach in building the current models. In addition to the reasons listed above, we believe that a fundamental approach was appropriate because the models used in Simulations 1 and 2 are extensions/variations on larger-scale models from previous studies (Harm & Seidenberg, 1999; Plaut et al., 1996). At the same time, we acknowledged the virtues of the realist approach by trying to include any potentially relevant details, so long as they did not interfere with our ability to implement the core principles.

Simulation 1

In Simulation 1, we instantiated the input gain hypothesis in a specific model of word reading based on the “triangle” framework (Plaut et al., 1996; Seidenberg & McClelland, 1989), so-called because it is often diagrammed in the shape of a triangle (see Fig. 4). Previous studies have implemented various approximations of the triangle framework (e.g., Harm, 1998; Harm & Seidenberg, 1999; Plaut, 1997; Plaut et al., 1996) but none of these implementations were equipped to model results from tempo-naming or stimulus blocking experiments. Therefore, we re-implemented the triangle framework in Simulation 1. Our re-implementation retained many details of the computational methods used in previous implementations, but some methods were altered to simulate the tempo-naming and stimulus blocking experiments. All computational details were motivated by a core set of theoretical principles that are outlined throughout the remainder of this section.

The model architecture used in the current simulation was based on three distinguishing characteristics of the triangle framework. First, the triangle framework

has been cast in terms of a distributed, connectionist network with three stipulated levels of representation: orthography, phonology and semantics. Second, orthographic representations are mapped to phonological representations through a semantic pathway and a phonological pathway (the mapping between semantics and phonology is learned prior to the introduction of orthography; see Methods). Third, the mappings in both pathways are learned in tandem with the task of minimizing error between actual and targeted semantic and phonological outputs. Coincident learning of the two pathways served to coordinate their processing (i.e., learning is interactive), thereby giving rise to a division of labor (see Harm, 1998; Plaut, 1997; Plaut et al., 1996).

In both pathways, learned, internal (i.e., hidden) representations mediate the mapping from orthography to phonology. However, in the semantic pathway, one of the intermediate layers of representation is explicitly semantic in structure (i.e., not hidden). Consequently, semantic structure is forced on the semantic pathway. By contrast, the phonological pathway is mediated only by hidden representations. Due to the nature of back-propagation learning, these hidden representations come to “blend” the structure in their inputs (orthography) and outputs (phonology; see Plaut & Gonnerman, 2000; Rumelhart, Durbin, Golden, & Chauvin, 1995).

In addition to theoretical principles that underlie the triangle framework, the current simulation was also guided by our desire to test the rate of processing hypothesis. This desire led to three particular modeling specifications: (1) the parameter input gain needed to be implemented in a model of word reading with an explicit time course of processing, (2) naming durations needed to be simulated, in addition to naming latencies and errors, and (3) the semantic pathway needed to be explicitly modeled. This last specification arose because the individual contributions of lexical and sub-lexical knowledge were key aspects of the tempo-naming results. The semantic pathway is important in this context because, in the triangle framework, the semantic pathway captures primarily lexical knowledge, whereas the phonological pathway captures primarily sub-lexical knowledge.

These three specifications were implemented as follows: (1) during processing, unit activations propagated across connections in continuous time, (2) two activation thresholds were placed on phonological outputs, one to represent the onset of a naming response, and the other to represent its offset, and (3) semantic knowledge was implemented as a high-dimensional space with hierarchical clusters of similarity among items, each cluster abstractly representing a “semantic” category (for a similar approach, see Plaut, 1995b; Plaut & Booth, 2000).

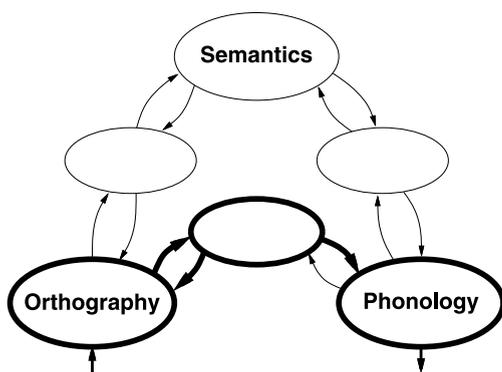


Fig. 4. A schematic of the “triangle” framework of word reading, adapted from Seidenberg and McClelland (1989).

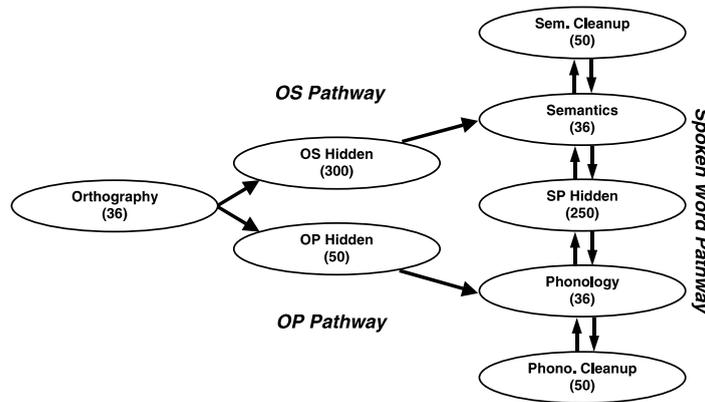


Fig. 5. Overview of the triangle architecture used in Simulation 1. The numbers in parentheses indicate the number of units in each group. Arrows indicate full connectivity. The written word pathway is comprised of the OS and OP pathways.

Methods

Network architecture. The network architecture is depicted in Fig. 5. Rather than divide the architecture in terms of a semantic and a phonological pathway, it is useful to divide it into into a *spoken word* pathway and two *written word* pathways. The spoken word pathway consisted of a single internal layer of representation, the *SP hidden layer*, that mediated the mappings to and from semantics and phonology. The written word pathways connected orthography with the spoken word pathway via the *orthography-to-semantics* (OS) pathway and the *orthography-to-phonology* (OP) pathway.³ The OS pathway mapped onto semantics via the *OS hidden layer*, whereas the OP pathway mapped onto phonology via the *OP hidden layer*. The number of hidden units in the OS pathway was much greater than that in the OP pathway (300 compared with 50, respectively). This difference was due to the fact that, in English, the mapping from orthography to phonology is much more systematic than the mapping from orthography to semantics. Systematic mappings are computationally more simple and therefore require fewer hidden units compared with arbitrary mappings.

Semantics and phonology were each bi-directionally connected to a group of “clean-up” units. Each clean-up group was trained to learn the structure of its inputs (see Training procedure). The written word pathways were made unidirectional for the sake of simplicity and because none of the core issues of the current work are dependent on interactions with orthography. These pathways would be bidirectional in a more complete version of the triangle framework.

Training corpus. A sample of 470 words was taken from a corpus of 2802 English monosyllabic, monomorphemic words. A sample was used as opposed to the full corpus to reduce the amount of computing resources necessary for training and testing. Words were sampled explicitly to preserve the distributional characteristics of the full corpus. It has been repeatedly demonstrated that human behavior is sensitive to the statistical properties of the environment. Given this fact, the distributions of words and their components are likely to impinge upon the processing of those words. Because connectionist models are also sensitive to the statistical properties of their inputs, we tried to preserve the distributions of monosyllabic word forms in English.

Sampling was accomplished by choosing the 12 most frequent onsets, vowel clusters, and codas in the full corpus.⁴ All words from the full corpus that could be made with these units were included in the sample corpus. This sampling method preserved the the pattern of frequency distributions of both orthographic and phonological onsets, vowel clusters, and codas in the full corpus. To preserve the distribution of word frequencies in the full corpus, sampled words were rank-ordered by their Kucera and Francis (1967) printed frequency. Frequencies were then distributed along the ordered list of words according to Zipf’s law by using the formula $F_w = 1/\sqrt{r + 2}$, where F_w is the frequency weighting for word w and r is that word’s rank order. This method modeled the general shape of the distribution of word frequencies found in English corpora, as well as corpora from other languages.

In addition to preserving frequency distributions, we wanted to preserve the lexical density of the full corpus. Lexical density can be defined as the proportion of words relative to the full size of the orthographic or

³ The OP pathway is synonymous with the phonological pathway, and the combination of the OS pathway with the spoken word pathway is synonymous with the semantic pathway.

⁴ A vowel cluster was defined as all vowel letters and the consonant “r” when it was adjacent and to the right of a vowel.

phonological space. Densities for English have been estimated to be in the range of 30 to 40% (Dell & Reich, 1981; Garrett, 1976). For the sample corpus, orthographic and phonological spaces were defined as all possible combinations of onsets, vowel clusters, and codas. By this definition, both the orthographic and phonological density of the sample corpus was 27%.

Finally, we wanted to preserve the distribution of mappings from spelling to sound found in the full corpus. We quantified consistency in terms of the number of alternate pronunciations for each orthographic onset, vowel cluster, and coda found in the full corpus, compared to that found in the sampled corpus (analogous to the measure of body consistency, e.g., Jared & Seidenberg, 1990). The distributions of alternate pronunciations in the sampled corpus matched the overall pattern of distributions in the full corpus.

Orthographic and phonological representations. Orthographic and phonological representations were based on those used in Plaut et al. (1996). Each monosyllabic word was represented by one onset, one vowel cluster, and one coda. Each orthographic and phonological processing unit represented a single onset, vowel cluster, or coda.⁵ In this coding scheme, units of a particular type (e.g., orthographic onsets or phonological codas) are pitted in competition with each other. To explicitly implement a mechanism of competition, unit activations were normalized to one over each of the six different types of units (i.e., softmax units).

For orthography, there were 12 onsets, 12 vowel clusters, and 12 codas. Pronunciations of the 470 words in the training corpus dictated that, for phonology, 12 onsets, 13 vowels, and 11 codas were necessary to represent all of the words.

Semantic representations. Semantic representations consisted of abstract categories with superordinate, basic, and subordinate category structure. There were four superordinate categories, 12 basic categories per superordinate category, and 16 subordinate features that were shared among the categories (although each feature was more likely to be activated for some categories, and less likely for others). Superordinate categories were represented by activating 2 out of 4 semantic units, basic categories were represented by activating 2 out of 16 semantic units, and subordinate features were represented by activating 4 out of 16 semantic units. Thus, 8 out of a total 36 semantic units were activated for each word. A semantic representation was generated for each word by randomly choosing a superordinate and basic category, and then probabilistically activating a subset of subordinate features. The semantic representations

for every possible pair of words differed by at least one unit (11.2 units on average).

The dimensions of semantic space were abstract and the placement of words in the semantic space bore no relation to their actual meanings. Our assumption was that the structure of semantic space may play a role in understanding the effect of pressure for speed on the word reading system, but the alignment of that structure with orthography and phonology was important in only one respect: apart from morphology, the components of orthography and phonology bear no systematic relationship with semantics.

Hidden representations. Hidden unit activations were constrained in the range $[-1:1]$ using the hyperbolic tangent output function, $\tanh x$. A polarized output function such as the hyperbolic tangent encourages positive feedback in recurrent networks, which facilitates point attractors at the targeted corners of unit space. A second property of the hyperbolic is that a net input of zero produces an output of zero. Thus, zero net input naturally corresponds to a state of no information.

Training procedure. There were two successive phases of training. In Phase 1, the spoken word pathway was learned without input from orthography. After the spoken word pathway was learned to criterion, all weights in the spoken word pathway were frozen. Phase 2 followed in which the written word pathways were learned.

This two-staged aspect of training embodied the natural time course of spoken language acquisition and reading acquisition. The relationship between the spoken and semantic forms of many words is learned well before their written forms. This fact was captured in the current simulation mainly by training on the mapping between semantics and phonology prior to training on the mappings from orthography. In addition, learning was halted in the spoken word pathway once the orthographic pathways were introduced. This latter feature may seem unrealistic if the acquisition of written word forms is considered to alter the spoken word pathway. However, we would argue that reading acquisition does not alter the spoken language system in any fundamental sense.⁶

Phases 1 and 2 shared a number of training methods in common. To reduce computational demands, training was carried out in feed-forward networks without a time

⁵ To represent words without onsets or codas, one unit represented the lack of an onset, and one unit represented the lack of a coda.

⁶ Because spoken language is still practiced during reading acquisition, one might imagine that a more realistic simulation would continue training on the spoken word pathway after the OS and OP pathways are introduced. This procedure should have an effect similar to the less time-consuming procedure of simply freezing the spoken word pathway (see McClelland, McNaughton, & O'Reilly, 1995 for simulations demonstrating how such "interleaved" learning prevents interference to previously learned knowledge).

course of processing (but the model was tested in continuous time; see Testing procedure). At the start of training, connection weights were initialized at random from a uniform distribution in the range ± 0.2 . At the beginning of each training example, unit outputs were initialized to 0. Error between outputs and external targets was scaled by word frequency, and output error was back-propagated to calculate weight derivatives using the formula

$$\Delta w_{ij}^{[t]} = \epsilon \left(\frac{\partial E}{\partial w_{ij}} + \alpha \Delta w_{ij}^{[t-1]} \right), \quad (4)$$

where t was the epoch number, the learning rate ϵ was set to 0.0005, and the momentum α was set to zero for the first 100 epochs, and 0.8 afterwards. Weights were updated after every 200 training examples.

Training: Phase 1. To train the bi-directional spoken word pathway in a feed-forward network, the pathway was unfolded such that semantics and phonology served as both the inputs and outputs to the SP hidden units (i.e., an auto-associator architecture). To train the bi-directional connections with clean-up units, semantic and phonological outputs generated from the SP hidden group projected into their respective clean-up groups, and the clean-up groups projected to a second layer of semantic and phonological outputs. Targets were set at both layers of output, but weight derivatives were not back-propagated across the first layer of output. Training a feed-forward network in this way is analogous to training a network using back-propagation through time for a single time step (Williams & Peng, 1990).

Three different training examples were created for each word. For semantics examples, semantic inputs were clamped, and phonological inputs were set to zero. For phonology examples, phonological inputs were clamped, and semantic inputs were set to zero. For full examples, both types of inputs were clamped. All three example types were pooled together, and training examples were chosen from the pool at random (with replacement). The net inputs to input units were calculated on the basis of their desired outputs according to their external input (i.e. soft-clamping). In addition, noise in the uniform range ± 0.5 was added to the net inputs of input units.

The spoken word pathway was trained on a total of three million examples, at which point the average error at the second layer of outputs was less than 0.01 per unit per example.

Training: Phase 2. In Phase 2, the written word pathways were trained in conjunction with the spoken word pathway, whose weights from Phase 1 were frozen. The OS pathway was connected to the semantic inputs of the spoken word pathway, and the OP pathway was connected to the phonological inputs. External inputs

were clamped to orthography, and noise in the uniform range ± 1.1 was added to the net inputs of orthography. There were 470 training examples, one for each word. Training procedures were similar to those of Phase 1 in other respects. The network was trained on a total of 10 million examples, at which point the average error at the second layer of outputs was less than 0.01 per unit per example.

Testing procedure. To test the network architecture shown in Fig. 5, the spoken word pathway was folded on itself such that respective input and output groups were merged. To allow unit outputs to propagate through the network in time, net inputs were time averaged as in Eq. (2). The time constant τ was set to 0.2. For each time step of a given test example, orthography was clamped with the given external inputs, and noise in the uniform range of ± 0.2 was added to the net inputs of all units in the network.

Orthographic test inputs included all training words and 814 nonwords. Nonwords included all possible orthographic onset–body (i.e., vowel plus coda) combinations given the training words. Onset–body nonwords were tested because it is a common procedure in empirical studies to create nonwords by shuffling the orthographic onsets and bodies of words (as was done in Kello & Plaut, 2000).

The network's response to each orthographic test input was measured at phonology. Naming latency corresponded to the time step at which the unit outputs of a phonological onset, vowel cluster, and coda exceeded 0.4. Naming offset corresponded to the time step at which the output of one unit of each type exceeded 0.95 (with a maximum of 40 time steps from the start of a given test example). Naming duration was measured as the time difference between naming offset and naming latency. This measure of naming duration is, admittedly, very coarse. A better approach would be to model the articulatory process in more detail, and in fact, work on such a model is in progress (Plaut & Kello, 1999). However, whole word duration is a coarse measure of response execution, and we would argue that our simulated measure of duration is adequate for this purpose. Our argument would be supported to the extent that the current simulation can account for the effect of tempo on naming durations.

The naming response corresponded to the most active phonological onset, vowel cluster, and coda during the processing of a given orthographic input. The network was tested on all training words and all defined nonwords (i.e., all novel onset–body combinations). If the naming response did not correspond to the targets at phonology, then it was placed in one of four different error categories: *LARC*, *word*, *mixed*, or *nonword*. For nonwords, targets at phonology included all legitimate mappings from orthography to phonology, i.e. all mappings that occurred in the training corpus. LARC

errors were legitimate, but incorrect phonological outputs given the orthographic input (LARC errors were not possible for nonwords). Word errors were incorrect outputs that formed one of the training words. Mixed errors were LARC errors that also formed words. All other errors were categorized as nonword errors. In the tempo-naming studies, errors included a fifth category, labeled “articulatory,” in which responses that were garbled beyond recognition were placed. To simulate articulatory errors, the process of articulation would need to be modeled in much more detail than implemented in the current simulation.

To acquire stable estimates of network performance, the network was tested ten times per test input, and results were averaged across all ten test runs. Errors were removed from analyses of latency and duration.

Standard naming results

The main purpose of the current simulation was to test the rate of processing hypothesis, but given that the model is a new implementation of the triangle framework, it is important to test the model against certain hallmark findings from experiments with the standard naming task. We chose to test the model against the well-known effects of printed frequency, spelling–sound consistency, and lexicality. In particular, it has generally been found that responses to high frequency words are faster and less error prone than those to low frequency words. Also, responses to words with exceptional SSCs (*exception words*) are slower than those to without such correspondences (*consistent words*). These two variables interact such that responses to low frequency exception words are slower and more error prone than those to the other three possible types of words. Finally, responses to pronounceable nonwords are generally slower and less accurate than those to words, although nonword performance can be comparable to performance for difficult words. Nonwords also served to test the model on its ability to generalize, which is another hallmark issue in the word reading literature (for a review of these effects and issues, see Plaut et al., 1996; Seidenberg & McClelland, 1989).

Words from the training corpus were chosen according to criteria along the dimensions of frequency

and spelling–sound consistency. High frequency words had frequencies greater than 0.15, and low frequency words had frequencies less than 0.07. Consistent words had spelling–sound mappings that were all greater than 70% consistent, exception words had at least one spelling–sound mapping that was less than 50% consistent. These dimensions were crossed to create four categories of words (numbers of each type shown in parentheses): high frequency consistent (24 HFC words), high frequency exception (15 HFE words), low frequency consistent (232 LFC words), and low frequency exception (25 LFE words).

The network was tested on these stimuli at an input gain of 1.0. Latencies and error percentages are shown in Table 1 as a function of frequency, consistency, and lexicality. Simulation results are shown along side the standard naming results in Experiment 2 from Kello and Plaut (2000). Both model and empirical results show that LFE words are slower and more error prone than either HFE or LFC words. Model results differed slightly from the empirical results in that nonwords latencies were faster than LFE words, and LFC latencies were faster than HFC latencies.

For latencies, the main effects of frequency and regularity were reliable (all ANOVAs are reported by items), $F(1, 290) = 12.0$, $p < .001$ and $F(1, 290) = 23.6$, $p < .001$, respectively. Although the interaction pattern was as predicted (a stronger effect of regularity for LF compared with HF words), the interaction effect was not reliable, $F(1, 290) = 1$. For error rates, all three effects were reliable, $F(1, 292) = 9.5$, $p < .01$, $F(1, 292) = 7.1$, $p < .01$, and $F(1, 292) = 7.1$, $p < .01$, respectively. Results show that, in a simulation of standard naming, the network exhibited the hallmark effects of frequency and regularity (albeit the power by items was apparently too low to show a reliable interaction effect for latencies). Moreover, training on words generalized to nonword performance. These results are sufficient to validate our implementation of the triangle framework.

Tempo-naming results

Pressure for speeded responding, as caused by fast tempos in the tempo-naming task, was hypothesized to

Table 1
Standard naming latencies for Sim. 1 and Expt. 2 from Kello and Plaut (2000) with error percentages shown in parentheses

	Simulation			Empirical		
	HF	LF	Nwd	HF	LF	Nwd
Con	8.43 (0.0)	8.92 (0.0)		NA (NA)	473 (3.5)	
Exc	9.25 (0.0)	10.15 (11.2)		471 (3.4)	488 (16.3)	
Nwd			9.75 (2.4)			510 (8.5)

HF, high frequency; LF, low frequency; Nwd, nonword; Con, consistent; Exc, exception; NA, not available.

cause an increase in the rate of processing in the word reading system. To simulate an increase in rate of processing, input gain was increased uniformly across all units in the network. Levels of input gain were chosen such that the overall error percentages were comparable to those found in Kello and Plaut (2000). To best match the stimulus conditions from Kello and Plaut (2000), test inputs were chosen to fit the four categories used in that study: HFE, LFE, LFC, and nonwords. In computing means as a function of input gain, each category was weighted evenly to match the specific conditions of Experiment 2 from Kello and Plaut (2000). This weighting compensated for the fact that each category had a different number of items. The conditions of Experiment 2 were chosen because this experiment contained all stimulus types used in that study. Moreover, findings were essentially the same for all three experiments in the tempo-naming study.

Naming behavior as a function of input gain is shown in Fig. 6, with the tempo-naming results from Experiment 2 of Kello and Plaut (2000) shown on the right. In the top-left panel, naming latencies and durations can be seen to decrease as input gain increases, with smaller decrements at higher levels of input gain. In the bottom panel, percentages for all types of errors can be seen to increase as input gain increases. The figure shows that, at higher levels of input gain, nonword

errors increase the most, followed by LARC errors, word errors, and mixed errors, respectively. To compare simulation and empirical results, the range of input gain values is delineated (the dashed lines) for which the overall model error rates were similar to those found empirically.

To test for the influence of input gain on stimulus effects, the effects of frequency, consistency, and lexicality (as defined for the standard naming analyses) were analyzed as a function of input gain. The size of all three stimulus effects decreased monotonically as input gain increased. This general result can be understood as a compression of stimulus effects that occurs in conjunction with a compression of the response trajectory. As discussed earlier, one result from the tempo-naming study was that stimulus effects were attenuated in the tempo-naming task compared with the standard naming task. Input gain partially accounts for this result due to the compression of stimulus effects. A full account would require the implementation of a timing mechanism, which is beyond the scope of the current study.

Blocking results

The mechanism of input gain was hypothesized to account for blocking effects such as those reported by

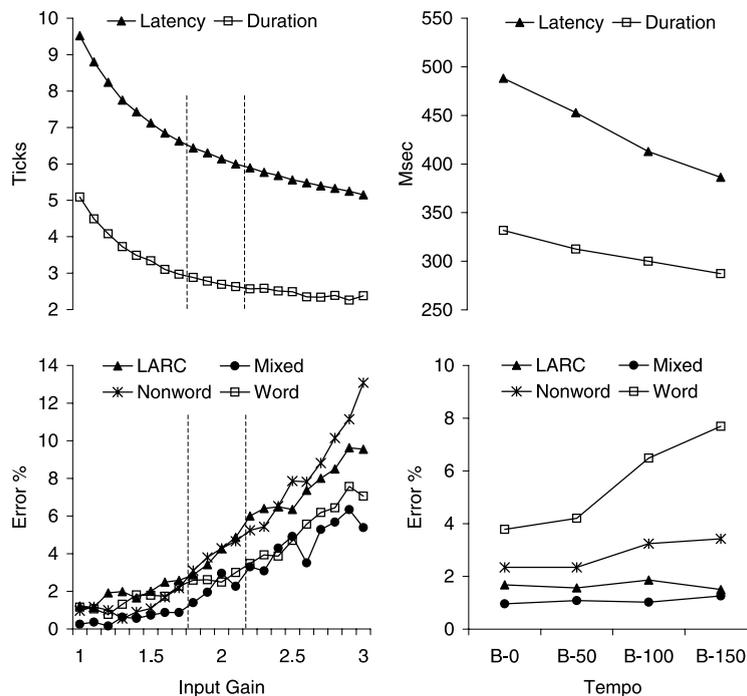


Fig. 6. Mean naming latencies, durations, and error percentages as a function of input gain in Sim. 1 and Expt. 2 from Kello and Plaut (2000). For the simulation, latencies and durations are given in numbers of ticks of processing. Articulatory errors in the empirical data were excluded for ease of comparison.

Jared (1997) and (Lupker et al. (1997)) in a fashion similar to the time criterion hypothesis. Difficult stimuli (i.e., more error prone) encourage lower levels of input gain, whereas easier stimuli encourage higher levels of input gain. When more and less difficult stimuli are mixed together, input gain is set at a middling level.

To test input gain as an alternative to a time criterion, we simulated Experiments 1, 2, and 3 from the study by Lupker et al. (1997). For pure blocks, we set input gain levels to match the rank orders of latencies and errors for the different types of stimuli tested in those experiments. As with a time criterion, we also assumed that levels of input gain are lower for more difficult stimuli to improve accuracy. For mixed blocks, we averaged the pure block levels of input gain. Results (given in Table 2) showed that the effect of input gain on latencies was analogous to that of a time criterion: levels of input gain were homogenized in mixed blocks compared with pure blocks, which cause latencies to become homogenized as well. This pattern of latency effects closely matched the effects reported in Lupker et al. (1997). In addition to affecting latencies, shifts in input gain caused small speed/accuracy tradeoffs in many of the comparisons between pure and mixed blocks. Empirical results were characterized by speed/accuracy tradeoffs in some, but not all, comparisons. However, error effects were marginal in both the simulation and empirical results.

Simulation 1: Discussion

The main purpose of Simulation 1 was to test the rate of processing hypothesis as an account of the tempo-

naming results, and as an alternative to the time criterion hypothesis. The simulation showed that input gain can mostly account for the effect of tempo: higher levels of input gain caused faster, shorter, more erroneous responses, and smaller stimulus effects. These results confirm the analyses reported in “Expected effects of input gain.” The simulation showed also that input gain can account for stimulus blocking results in much the same way as a time criterion. Furthermore, our implementation of the triangle framework was validated by the demonstrated effects of frequency, consistency, and lexicality.

Despite the overall fit between Simulation 1 and empirical findings, there were two mismatches that may be important for understanding the underlying mechanisms. First, the blocking results of Simulation 1 produced longer latencies to LFE words compared with nonwords in the mixed block. Lupker et al. (1997) found the opposite result in Experiment 1 of that study, and the authors explained that the time criterion hypothesis makes the same incorrect prediction as demonstrated for input gain in Simulation 1. To address this shortcoming, Lupker and his colleagues proposed a “lexical-checking” mechanism in which readers can strategically choose to compare phonological outputs against an output lexicon. They proposed that lexical-checking is active for pure blocks of LFE words, non-active for pure blocks of nonwords, and partially active in the mixed block. Given that lexical-checking would tax the system and thereby increase latencies, the mechanism would account for the result.

The failure of input gain to account for the result in Simulation 1 may also be explained by the addition of a lexical-checking mechanism. However, a second ex-

Table 2
Blocking results for Simulation 1

Stimulus type	Simulation results						Lupker et al.'s results					
	Pure		Mixed		Effect		Pure		Mixed		Effect	
	RT	ER	RT	ER	RT	ER	RT	ER	RT	ER	RT	ER
<i>Expt. 1</i>												
HFE	7.1	1.3	7.5	0.7	+0.4	-0.6	463	2.3	485	2.7	+22	+0.4
LFE	9.3	11.6	8.7	11.6	-0.6	0.0	563	9.4	547	11.8	-16	+2.4
Nwd with HFE	7.6	3.6	7.3	3.9	-0.3	+0.3	554	4.8	535	8.2	-19	+3.4
Nwd with LFE	7.6	3.6	8.1	3.1	+0.5	-0.5	555	4.4	564	7.0	+9	+2.6
<i>Expt. 2</i>												
HFC	6.6	0.0	6.8	0.7	+0.2	+0.7	442	1.4	469	1.4	+27	0.0
LFC	6.9	2.7	7.1	2.1	+0.2	-0.6	482	2.2	496	2.9	+14	+0.7
Nwd with HFC	7.6	3.6	7.3	3.9	-0.3	+0.3	541	6.2	520	9.8	-21	+3.6
Nwd with LFC	7.6	3.6	7.3	3.9	-0.3	+0.3	552	6.6	532	6.9	-20	+0.3
<i>Expt. 3</i>												
HFE	7.7	0.7	8.1	0.0	+0.4	-0.7	488	3.6	513	1.8	+25	-1.8
LFE	9.3	11.6	8.7	11.6	-0.6	0.0	583	12.0	559	12.2	-24	+0.2

Simulated experiments are from Lupker et al. (1997).

planation is possible. Input gain may have failed to account for this result because, at any given level of input gain, latencies to LFE words were faster than those to nonwords. Therefore, when input gain was homogenized in the mixed block condition, latencies to LFE words became faster than those to nonwords. However, if simulated latencies to LFE words were slower in general than those to nonwords, then latencies in a simulated mixed block would match the pattern observed by Lupker and his colleagues. This alternative explanation does not require any additional mechanisms. The explanation is made possible by the fact that the effect of input gain on latencies interacts with the dynamics of the word reading system. By contrast, the time criterion operates independently of such dynamics.

The second mismatch between Simulation 1 and empirical findings was in simulating the tempo-naming task. As input gain was increased, there was a large increase in LARC and mixed errors, and LARC errors occurred more frequently than word errors. The tempo-naming results showed a different pattern: as the tempo was increased, there was little or no increase in LARC or mixed errors, and word errors occurred more frequently than LARC errors. To understand this discrepancy, we need to consider the role of SSCs in mapping from orthography to phonology. In the “Expected effects of input gain,” we pointed out that the occurrence of LARC and mixed errors (above chance) presumably arises from the erroneous influence of SSCs. Thus, to account for the pattern of errors found in the tempo-naming study, an increase in input gain would need to attenuate the influence of SSCs on the computation of phonology.

In the triangle framework, SSCs are represented primarily within the OP pathway, and lexical-sized correspondences are represented primarily within the OS pathway. This division of labor suggests that, in Simulation 1, the influence of the OP pathway on the computation of phonology was too strong, relative to that of the OS pathway. The strength of the phonological pathway is not surprising if one considers that in English, the relationship between orthography and phonology is much more systematic than that between orthography and semantics (Plaut et al., 1996; Van Orden & Goldinger, 1994). Internal representations are much easier to learn within connectionist networks if similar inputs tend to correspond to similar outputs (Rumelhart et al., 1995).

How can the influence of the OP pathway be reduced in the triangle framework? One possibility would be to reduce the systematicity between orthography and phonology, or increase the systematicity between semantics, orthography, and phonology. These manipulations did not seem warranted to us for two reasons. First, we created the training corpus to preserve syste-

maticity and other distributional characteristics of the full English corpus. Furthermore, the implementation reported by Plaut et al. (1996) was shown to exhibit an error pattern similar to that of Simulation 1 (results reported in Kello & Plaut, 2000).

A second possibility would be to either reduce the level of input gain in the OP pathway, or increase the level of input gain in the OS pathway. In doing so, input gain might serve as a mechanism analogous to the route emphasis hypothesis: the contribution of each pathway to phonological activation would be a function of the proportional difference in levels of input gain between the two pathways. Such a use of input gain would be analogous to how input gain was used by Cohen and his colleagues to modulate the influence of specific sources of information on processing (Cohen et al., 1990; Cohen & Servan-Schreiber, 1992). On this account, the predominance of word errors observed in the tempo-naming experiments would be explained as increased emphasis of the OS pathway as a function of tempo. This explanation was tested in Simulation 2.

Simulation 2

In Simulation 2, input gain was manipulated separately in the OS and OP pathways as a mechanism of route emphasis in the triangle framework (for another computational implementation of route emphasis, see Zorzi, 1999). We hypothesized that variable levels of input gain in these two pathways would alter the pattern of errors as compared with the pattern found in Simulation 1. In particular, a higher level of input gain in the OS pathway compared with the OP pathway was hypothesized to provide a closer simulation of the tempo-naming results. However, to provide a more complete test of input gain as a mechanism of route emphasis, higher levels in both pathways were tested.

Methods

The network and testing procedures from Simulation 1 were used in Simulation 2. Emphasis on the OP pathway corresponded to an increased level of input gain, γ , for OP hidden units, and a decreased level of input gain, $1/\gamma$, in OS pathway and the spoken word pathway (which comprise the semantic pathway). Emphasis on the OS pathway corresponded to the inverse application of input gain.⁷ Input gain was in-

⁷ A less extreme method of emphasis would be to increase or decrease input gain in only one pathway, rather than the complimentary method that we used. In an unreported simulation, we implemented the less extreme method and found that it was too weak in terms of the amount of emphasis that it placed on lexical versus sub-lexical contributions to processing.

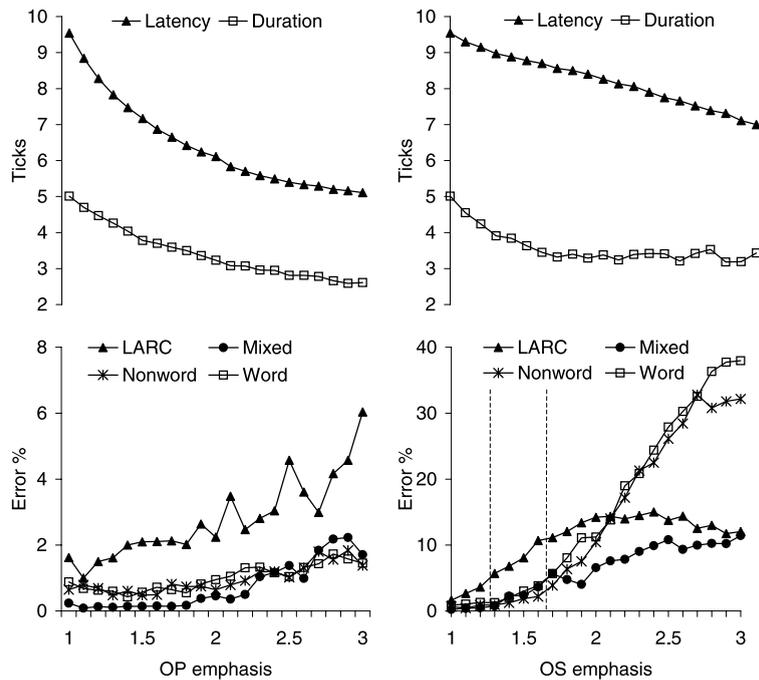


Fig. 7. Mean naming latencies, durations, and error percentages as a function of emphasis in Simulation 2.

creased for phonological units and phonology cleanup units for both types of emphasis. Emphasis always corresponded to increased input gain at phonology because, in order to simulate the tempo-naming task, emphasis must be associated with shorter latencies and durations. Under the rate of processing hypothesis, shorter latencies and durations are caused by increased levels of input gain, particularly at phonology.

Tempo-naming results

As in Simulation 1, stimuli were chosen to match the conditions from Experiment 2 of Kello and Plaut (2000). Naming behavior as a function of emphasis is shown in Fig. 7. In the top panels, naming latencies and durations can be seen to decrease overall as emphasis increases, with smaller decrements at higher levels of emphasis. For emphasis on the OS pathway, naming durations cease to increase at moderately high levels of emphasis. In the bottom panels, percentages for all types of errors can be seen to increase overall with higher levels of emphasis. As in Fig. 6, the dashed lines delineate the range of input gain values for which the overall model error rates were similar to those found empirically. These lines are shown only for OS emphasis because this is the manipulation that simulated the effect of tempo.

Two large differences can be seen in the effects of OS versus OP emphasis on errors. First, emphasis on the OS pathway disrupted processing much more than emphasis

on the OP pathway. Second, emphasis on the OS pathway caused a much larger increase in word and nonword errors, compared with LARC and mixed errors. By contrast, LARC errors were predominant when emphasis was increased on the OP pathway.

Simulation 2: Discussion

The results of Simulation 2 showed that the distribution of naming errors across different error categories can be altered by route emphasis. When input gain was proportionally greater in the OP pathway compared with the OS pathway, LARC errors became predominant. Conversely, when input gain was proportionally greater in the OS pathway, word errors became predominant, but only at high levels of input gain.

The simulated pattern of errors is generally supportive of use of input gain as a mechanism of route emphasis, but a closer examination revealed a significant mismatches between the model and empirical findings. Error proportions from the tempo-naming study were matched only at high levels of OS emphasis. Within the delineated range of input gain values, the error patterns were very similar to those found in Simulation 1. Therefore, the error patterns in Simulation 2 do not match the empirical data.

Independent of the match between model and empirical results, one must ask whether emphasis on the OS pathway makes sense given the conditions of the tempo-

naming study. The OS pathway should be emphasized in conditions that favor lexical as opposed to sub-lexical knowledge. However, as explained in Kello and Plaut (2000), the conditions of Experiments 2 and 3 in that study did not seem to favor lexical knowledge because the stimuli included nonwords (25 and 100% nonwords, respectively). The OS pathway should interfere with nonword processing or, at best, be useless for nonword processing, because nonwords have no clear semantic link between orthography and phonology. Therefore, it would be puzzling if the OS pathway was emphasized to satisfy the demand for speeded responding.

The questionable basis of the route emphasis account led us to consider an alternative. LARC errors are caused primarily by the contribution of the OP pathway because the OP pathway must capture the sub-lexical SSCs during reading acquisition. The role of the OP pathway could be seen in the predominance of LARC errors when the OP pathway was emphasized in Simulation 2, and in the low proportion of LARC errors when the OS pathway was emphasized, relative to word errors. How might the contribution of the OP pathway be attenuated without using a mechanism of route emphasis? One possible solution was tested in Simulation 3.

Simulation 3

One reason for the strong influence of SSCs in the triangle framework may be that processing (as opposed to learning) in the OP pathway is mostly independent of processing the OS pathway. This independence allows the strong systematicity between orthography and phonology in English to override the semantically-mediated contribution of the OS pathway. If the processing of SSCs and semantic correspondences was more integrated, then the systematicity between orthography and phonology may play a lesser role due to close interactions with semantically-mediated processes.

To integrate the processing of SSCs and semantic correspondences, we altered the triangle architecture by mapping orthography onto the SP hidden layer (i.e., the *OSP pathway*), rather than onto semantics and phonology directly (see Fig. 8). This form of integration between the OS and OP pathways was motivated by certain characteristics of the relationship between reading and spoken language processing (for similar approaches, see Perfetti & Sandak, 2000; Zorzi, Perry, Ziegler, & Coltheart, 1999).

Spoken language skills are typically acquired to a high degree of proficiency prior to reading acquisition. In distributed connectionist models of word reading, the spoken language system is represented by the spoken word pathway. The precedence of spoken word acquisition was captured in Simulations 1 and 2 by training the spoken word pathway prior to training the written word pathways. Moreover, the written word pathways were forced to learn in the context of the spoken word pathway. Nonetheless, the influence of the spoken word pathway on learning in the written word pathways was weak and indirect; the error incurred at semantics was primary in shaping the OS hidden representations, and the error incurred at phonology was primary in shaping the OP hidden representations. Derivatives collected at the SP hidden layer did not impinge substantially on learning in the OS hidden and OP hidden layers.

The consequence of the triangle architecture is that the written word pathways do not take full advantage of the learning that occurs during the acquisition of spoken words. In learning the spoken word pathway, SP hidden representations are shaped to maintain the activation of both semantic and phonological representations. To simulate word reading, the network must map orthography onto both semantics and phonology. The function and properties of the SP hidden layer seem to present an opportunity for learning a mapping from orthography to phonology. In an *integrated-pathway* architecture, the reading system would form one pathway from orthog-

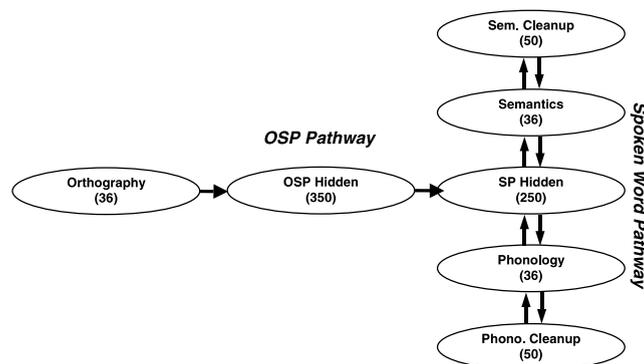


Fig. 8. Overview of the integrated-pathway architecture used in Simulation 3. The numbers in parentheses indicate the number of units in each group. Arrows indicate full connectivity. The written word pathway is synonymous with the OSP pathway.

raphy directly into the SP hidden layer (the OSP pathway). This *OSP pathway* could then support the activation of both semantic and phonological representations, which is typically the core task demand that is simulated in models of word reading. By contrast, the triangle framework requires two pathways to correctly activate both semantic and phonological representations (for a similar argument, see Frost, 1998).

The integrated-pathway architecture has an advantage over the triangle framework also in terms of its use of morphological representations. In a broader implementation of the spoken word pathway, the task demands that impinge upon the spoken word pathway would shape it to capture the morphological structure of words (Plaut & Gonnerman, 2000). In the triangle framework, morphological structure would similarly be learned in the OS pathway because the relationship between orthography and semantics is similar to that between phonology and semantics. Therefore, much of the morphological structure mediating the spoken word pathway would essentially be re-learned and re-represented in the OS pathway. By contrast, the morphological structure learned during spoken word acquisition is “reused” by the OSP pathway.

The argument for a OSP pathway may seem appealing, but if the OP pathway is removed, then it may seem difficult to account for nonword reading. The spoken word pathway is lexical by nature; words, but not nonwords, have explicit relationships between sound and meaning. Therefore, during the acquisition of spoken words, nonwords are not processed (or only partially processed) via the spoken word pathway. How, then, might nonwords be pronounced in a system that has only the OSP pathway to activate phonological representations from orthographic input?

The potential answer lies in the nature of distributed representations, and in particular, representations that are learned via error-correcting learning procedures like back-propagation. Hidden representations blend the structures of their inputs and outputs. Consequently, the SP hidden layer should be a blend of semantic and phonological structure. On a distributed connectionist approach, the ability to read nonwords emerges from the componential relationship between orthography and phonology. To the extent that the SP hidden layer is shaped by phonological structure, an OSP pathway

should be able to capitalize on the componential relationship between orthography and phonology. As a test of this logic, the ability to process nonwords was measured in the current simulation, as it was in Simulation 1.

Methods

The training corpus, representations, training procedures, and testing procedures used in Simulation 1 were also used in Simulation 3, with the following exceptions. In Phase 2 of training, orthography was mapped onto the SP hidden units via a single group of 350 hidden units (equal to the combined number of OS and OP hidden units in Simulation (1)). The network was trained in Phase 2 on a total of seven million examples.

Standard naming results

Latencies and error percentages are shown in Table 3 as a function of frequency, consistency, and lexicality. Simulation results are shown along side the standard naming results in Experiment 2 from (Kello & Plaut, 2000). Both model and empirical results show that LFE words are slower and more error prone than either HFE or LFC words. As in Simulation 1, model results differed slightly from the empirical results in that nonwords latencies were faster than LFE words, and LFC latencies were faster than HFC latencies.

For latencies, the main effects of frequency and regularity were reliable (all ANOVAs are reported by items), $F(1, 291) = 24.1$, $p < .001$ and $F(1, 291) = 118.8$, $p < .001$, respectively. Unlike Simulation 1, the interaction of these two variables was reliable, $F(1, 291) = 7.0$, $p < .01$. For error rates, all three effects were reliable, $F(1, 292) = 8.5$, $p < .01$, $F(1, 292) = 7.6$, $p < .01$, and $F(1, 292) = 7.6$, $p < .01$, respectively. As in Simulation 1, the network exhibited the hallmark effects of frequency and regularity. The network’s ability to use the systematicity between orthography and phonology in reading was evidenced in its performance on nonwords, albeit performance was not as high as in Simulation 1 (8.1% errors compared with 15.5% errors, respectively). These results provide initial validation for the integrated-pathway model (but see General discussion).

Table 3
Standard naming latencies for Sim. 3 and Expt. 2 from Kello and Plaut (2000), with error percentages shown in parentheses

	Simulation			Empirical		
	HF	LF	Nwd	HF	LF	Nwd
Con	7.93 (0.0)	8.20 (0.3)		NA (NA)	473 (3.5)	
Exc	8.92 (0.0)	9.86 (8.8)		471 (3.4)	488 (16.3)	
Nwd			9.0 (11.5)			510 (8.5)

Tempo-naming results

As in Simulations 1 and 2, stimuli were chosen to match the conditions from Experiment 2 of Kello & Plaut (2000). Naming behavior as a function of input gain is shown in Fig. 9. In the top panel, naming latencies and durations can be seen to decrease as input gain increases, with smaller decrements at higher levels of input gain. At extreme values of input gain, naming duration can be seen to increase slightly. In the bottom panel, percentages for all types of errors can be seen to increase as input gain increases. At higher levels of input gain, word and nonword errors increase the most, with LARC and mixed errors increasing less. As in Simulations 1 and 2, the dashed lines delineate the range of input gain values for which the overall model error rates were comparable to those found in Experiment 2 from (Kello & Plaut, 2000).

Blocking results

As an account of stimulus blocking, the manipulation of input gain should have the same effect in the integrated-pathway model as it did in the triangle model in Simulation 1. Difficult stimuli should encourage lower levels of input gain, whereas easier stimuli should encourage higher levels. When stimuli of greater and lesser difficulty are mixed together, input gain would be set at a middling level. Using the same method as described in Simulation 1, we tested the rate of processing account of stimulus block in the integrated-pathway model. Results are shown in Table 4. As in Simulation 1, the pattern of latency effects matched those found in (Lupker et al., 1997), and the pattern of error rates showed small speed/accuracy tradeoffs for each comparison.

Simulation 3: Discussion

Simulation 3 served three purposes: (1) it provided a second test of the rate of processing hypothesis, (2) it

tested a novel alternative to the triangle framework against hallmark phenomena in word reading, and (3) it tested whether the naming responses of an integrated-pathway architecture would better match the distribution of errors found in the tempo-naming study compared with our implementation of the triangle framework in Simulation 1.

With regard to the first purpose, input gain was shown to modulate naming latencies, durations, and overall error rates in a manner consistent with the tempo-naming results. In addition, input gain was shown to account for stimulus blocking results in manner analogous to a time criterion account. These results show that, as a mechanism of control over rate of processing, input gain had very similar effects in the integrated-pathway model compared with the triangle model. This consistency across two different computational architectures provides converging support for the rate of processing hypothesis.

With regard to the second purpose, the integrated-pathway model produced behavior in a simulation of the standard naming task that was very similar to that produced by the triangle framework in Simulation 1. The only notable difference was in performance on reading nonwords, and it is not surprising that the integrated-pathway model performed slightly worse than the triangle framework. The integrated-pathway framework was motivated in part by the need to reduce the influence of SSCs on naming behavior, which are responsible for the ability to read nonwords. In this light, the slight decrement in nonword performance is evidence that the integrated-pathway design was successful in reducing the influence of SSCs. Overall, the standard naming results in Simulation 3 provide support for the integrated-pathway architecture as a potentially viable model of word reading (although more work is necessary; see General discussion).

With regard to the third purpose, the integrated-pathway model produced a majority of word errors that increased with higher levels of input gain, whereas LARC and mixed errors increased less by comparison.

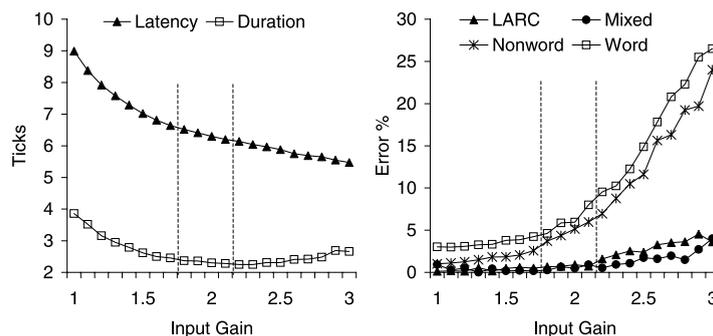


Fig. 9. Mean naming latencies, durations, and error percentages in Simulation 3 as a function of input gain.

Table 4
Blocking results for Simulation 3

Stimulus type	Simulation results					
	Pure		Mixed		Effect	
	RT	ER	RT	ER	RT	ER
<i>Expt. 1</i>						
HFE	6.9	0.7	7.3	0.0	+0.4	−0.7
LFE	9.1	6.4	8.5	7.6	−0.6	+1.2
Nwd with HFE	7.6	13.0	7.4	13.6	−0.2	+0.6
Nwd with LFE	7.6	13.0	8.0	12.2	+0.4	−0.8
<i>Expt. 2</i>						
HFC	6.5	0.0	6.7	0.0	+0.2	0.0
LFC	6.7	1.3	6.9	0.6	+0.2	−0.6
Nwd with HFC	7.6	13.0	7.4	13.6	−0.2	+0.6
Nwd with LFC	7.6	13.0	7.4	13.6	−0.2	+0.6
<i>Expt. 3</i>						
HFE	7.5	0.0	7.8	0.0	+0.3	0.0
LFE	9.1	6.4	8.5	7.6	−0.6	+1.2

Simulated experiments are from Lupker et al. (1997).

This general pattern matched the tempo-naming results more closely than the pattern produced in Simulation 3, in which LARC errors were in the majority. In light of all three simulations in this study, the error pattern produced by the integrated-pathway model was a function of both input gain and the integrated-pathway architecture. Results from Simulation 2 showed that high levels of input gain can diminish the contribution of SSCs even in an architecture in which SSCs had a relatively strong influence on processing. To complement, the contribution of SSCs was diminished in the current simulation, but not in Simulation 1. The only change from Simulation 1 to the current simulation was in the model architecture. Therefore, the difference in model behavior must have been due to the difference in architecture.

General discussion

Two experimental paradigms in word reading, tempo-naming and stimulus blocking, have shown that a mechanism exists to control the time course of processing in word reading. We hypothesized that rate of processing was controlled in response to the task conditions in the tempo-naming and stimulus blocking experiments, and we used input gain as a computational mechanism of control over rate of processing. Input gain was investigated in three connectionist simulations of word reading, and in all three simulations, input gain accounted for the overall pattern of latencies, durations, and errors found in the tempo-naming experiments. Moreover, in two different models of word reading, in-

put gain provided a consistent account of the stimulus blocking results reported in (Lupker et al., 1997). According to this account, difficult stimuli require low to moderate levels of input gain to preserve accuracy, whereas relatively easy stimuli allow for higher levels of input gain.

Despite the overall fit between empirical and computational results, there were two notable discrepancies. First, input gain did not account for a particular blocking result in which latencies to two types of stimuli of apparently equal difficulty (LFE words and nonwords) changed from pure to mixed block conditions. This failure can be explained by the fact that, in both reported models, latencies for LFE words were slower than those for nonwords. If latencies for LFE words were faster, then input gain would account for the result. Although this issue clearly requires further investigation, we should mention two possible ways in which simulated latencies to LFE words might be made faster than those to nonwords. First, in a more complete model of word reading trained on a full corpus of English words, latencies to all words (including LFE words) may be lower, relative to nonwords, compared with the current simulations. One might expect this to happen to the extent that lexical knowledge might have a stronger influence on processing when a model is trained on a much larger corpus. Second, if the actual words and nonwords from (Lupker et al., 1997) were used in the simulations, then latencies to LFE words may come out to be less than those to nonwords.

The second discrepancy was in the errors produced in Simulations 1 and 2 as a function of input gain, compared with errors produced in the tempo-naming ex-

periments as a function of tempo. In the tempo-naming experiments, word errors were most frequent and their frequency increased with tempo, whereas LARC and mixed errors were less frequent and these did not increase substantially with tempo. In Simulation 1, LARC and mixed errors occurred at a rate similar to word errors, and all types increased at higher levels of input gain. This mismatch in error patterns suggested that the influence of lexical knowledge was too weak relative to that of sub-lexical knowledge. This suggestion was supported in Simulation 2 in which word errors predominated over LARC and mixed errors when lexical knowledge was strongly emphasized by increasing input gain in the OS pathway and decreasing it in the OP pathway.

Issues with Simulation 2 led us in Simulation 3 to test an alternate method of manipulating the influence of lexical versus sub-lexical knowledge. In Simulation 3, the OS and OP pathways were combined into a single, integrated OSP pathway. Of all three simulations, the error pattern produced by the integrated-pathway model provided the closest match to the error patterns found in the tempo-naming study. The results of Simulation 3 also supported the hypothesis that lexical knowledge was too weak in Simulation 1, as well as the rate of processing hypothesis.

The current simulations raise a number of questions, two of which we address here. First, what are the implications of input gain for the two other proposed mechanisms of control in word reading? With regard to route emphasis, the results of Simulation 2 suggest that input gain can serve as a mechanism of route emphasis as well as a mechanism of control over rate of processing. A specific level of representation or processing can be emphasized by increasing input gain at that level, and/or decreasing it elsewhere in the system. Alternatively, rate of processing can be increased by increasing input gain across all levels of processing within that system. Although further investigation is necessary, the current simulations suggest that input gain may provide a unified account of two apparently different sources of strategic effects in word reading, one from studies supporting control over route emphasis (e.g., Zevin & Balota, 2000), and the other from studies supporting control over the time course of processing (e.g., Kello & Plaut, 2000).

With regard to a time criterion, the input gain mechanism seems to account for a broader range of phenomena (i.e., naming duration data and blocking data that are better accounted for by some mechanism of route emphasis). However, it is clear that input gain cannot be the sole mechanism of strategic control in word reading. For example, input gain does not explain how behavior is precisely entrained to an external rhythm in the tempo-naming task; it merely captures the hypothesis that rate of processing is increased in re-

sponse to faster tempos. To explain the entrainment of behavior to rhythms, a mechanism akin to a time criterion may be most appropriate (e.g., see Pashler, 2001). Further research is necessary to investigate how a rate of processing mechanism might interact with a mechanism akin to a time criterion to exert control over the word reading system.

A second question raised by the current simulations is whether the integrated-pathway model can stand as a general theory of word reading. The results of Simulation 3 provided initial support for the model by accounting for the well-known effects of frequency, regularity, and lexicality on standard naming latencies and errors. However, further computational work is clearly necessary to investigate this novel architecture. Findings from a variety of domains have been interpreted as evidence for a dual-pathway system of lexical processing. Such findings present a challenge for the integrated-pathway architecture because in this architecture, there is only processing pathway from orthography into the spoken language system. Perhaps the most challenging findings come from neuropsychological data that have been interpreted as support for a dual-pathway architecture. We briefly review these data here, and consider how they might be addressed in an integrated-pathway framework.

Two types of acquired dyslexia have received particular attention, in part because they seem to implicate a separation of lexical and sub-lexical processing mechanisms: surface and phonological dyslexia. Surface dyslexia (see Patterson, Coltheart, & Marshall, 1985) is characterized by a selective impairment in reading exception words, particularly those of low frequency, relative to regular words and nonwords; phonological dyslexia (see Beauvois & Derouesné, 1979; Coltheart, 1996b) is characterized by a selective impairment in reading nonwords relative to regular and exception words.⁸ Both surface and phonological dyslexia have been observed following brain injury in premorbidly literate adults (acquired dyslexia), and among children who failed to acquire age-appropriate reading skills despite adequate intelligence and education (developmental dyslexia; see Castles & Coltheart, 1993; Manis, Seidenberg, Doi, McBride-Chang, & Peterson, 1996).

The existence of these complementary disorders constitutes a double dissociation: the impairment of exception words coupled with the preservation of nonwords in surface dyslexia on the one hand, and the reverse condition in phonological dyslexia on the other.

⁸ A third major type of dyslexia, *deep* dyslexia (see Coltheart, Patterson, & Marshall, 1980), can be interpreted as a severe form of phonological dyslexia in which patients also make semantic errors (Friedman, 1996; Glosser & Friedman, 1990).

Double dissociations are often interpreted as evidence for the existence of anatomically and functionally distinct sub-systems or sub-processes (but see Plaut, 1995a; Van Orden, Haar, Jansen op de, & Bosman, 1997); in this case, separable lexical and sub-lexical reading sub-systems. In particular, within a dual-route theory, acquired surface and phonological dyslexia have natural interpretations as damage to the lexical and sub-lexical pathways, respectively (e.g., Coltheart, 1981; Coltheart, 1985; Coltheart et al., 1993); although see (Coltheart et al., 2001, for a different account).

In fact, an analogous account would seem to be available within the triangle framework, with damage to the semantic pathway causing surface dyslexia and damage to the phonological pathway causing phonological dyslexia. Indeed, there is extensive evidence for a close relationship between semantic damage and surface dyslexia (Balota & Ferraro, 1993; Graham, Hodges, & Patterson, 1994; Hillis & Caramazza, 1991; Patterson, Graham, & Hodges, 1994; Patterson & Hodges, 1992; although see Cipolotti & Warrington, 1995; Lambon Ralph, Ellis, & Franklin, 1995). According to Patterson and colleagues, semantic damage impairs performance on LFE words because these items are processed most weakly by the phonological pathway and come to rely on semantic support.

However, also following Patterson and colleagues (Patterson & Marcel, 1992; Patterson, Suzuki, & Wydell, 1996), Plaut et al. (1996) argued that phonological dyslexia is better understood as arising not from damage to the phonological pathway but from damage directly to phonology, in part because virtually all phonological dyslexic patients are impaired at purely phonological tasks with no reading component (see Coltheart, 1996b). This strong association is problematic for dual-route theories but also for the triangle framework (Coltheart, 1996), because both architectures permit damage to the sub-lexical/phonological pathway that spares phonology, which should produce phonological dyslexia without concomitant non-reading phonological impairments.

One possible response to this challenge within the triangle framework is that damage restricted to the mapping between orthography and phonology might, in fact, produce impairments in purely phonological tasks due to learned interactive support between orthography and phonology in literate individuals (Morais, Cary, Alegria, & Bertelson, 1979; Morais, Bertelson, Cary, & Alegria, 1986; Stone, Vanhoy, & Van Orden, 1997; Van Orden, Pennington, & Stone, 1990). Substantiating this proposal, however, would require a full implementation of the triangle framework capable of performing a broad range of both reading and non-reading tasks.

The integrated-pathway architecture provides an alternative perspective on these issues. Damage to the OSP pathway would be expected to produce a mixed pattern

of reading impairment across all stimulus types, which is in fact the most common pattern observed among dyslexics (see Manis et al., 1996). Damage to the SP hidden units within the spoken word pathway would produce similar reading impairments but now also impact spoken language performance. Only semantic damage (or the connections between semantics and the SP hidden units) should produce surface dyslexia by removing semantic support for LFE words, whereas only phonological damage should produce phonological dyslexia by impairing those items without semantic support. Like the triangle model, however, a full implementation of the integrated-pathway architecture remains necessary to substantiate this proposal.

Conclusion

The majority of past research on word reading has focused on automatic processes, such as stimulus and context effects. More recently, the issue of strategic control over reading processes has come into focus. In the current study, a mechanism of strategic control was investigated in three connectionist models of word reading. Simulation results supported the hypothesis that rate of processing is a control parameter over the word reading system. In future empirical studies, control parameters could be exploited to test the flexibility and boundary conditions of various cognitive processes. Such tests may provide insight into the general architecture of cognitive systems such as the ones that support reading skills.

Acknowledgments

The research was supported by a start-up grant awarded to the first author by George Mason University, and an NIMH FIRST award (MH55628) to the second author. The computational simulations were run using the Lens network simulator (version 2.3.3, modified for the simulations herein), written by Doug Rohde (<http://www.cs.cmu.edu/~dr/Lens>). We thank Marco Zorzi, Mike Harm, Alan Kawamoto, Jay McClelland, and the CMU PDP research group for helpful comments and discussion.

References

- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147–169.
- Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory & Cognition*, 10, 565–575.
- Balota, D., & Ferraro, R. (1993). A dissociation of frequency and regularity effects in pronunciation performance across

- young adults, older adults, and individuals with senile dementia of the Alzheimer type. *Journal of Memory and Language*, 32, 573–592.
- Baluch, B., & Besner, D. (1991). Visual word recognition: Evidence for strategic control of lexical and nonlexical routines in oral reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 644–652.
- Beauvois, M.-F., & Derouesné, J. (1979). Phonological alexia: Three dissociations. *Journal of Neurology, Neurosurgery, and Psychiatry*, 42, 1115–1124.
- Castles, A., & Coltheart, M. (1993). Varieties of developmental dyslexia. *Cognition*, 47, 149–180.
- Cipolotti, L., & Warrington, E. K. (1995). Semantic memory and reading abilities: A case report. *Journal of the International Neuropsychological Society*, 1, 104–110.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332–361.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, 45–77.
- Colombo, L., & Tabossi, P. (1992). Strategies and stress assignment: Evidence from a shallow orthography. In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology, and meaning* (pp. 319–399). Amsterdam: North-Holland.
- Coltheart, M. (1981). Disorders of reading and their implications for models of normal reading. *Visible Language*, 15, 245–286.
- Coltheart, M. (1985). Cognitive neuropsychology and the study of reading. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 3–37). Hillsdale, NJ: Erlbaum.
- Coltheart, M. (Ed.) (1996). Phonological dyslexia: Past and future issues. *Cognitive Neuropsychology*, 13, 749–762.
- Coltheart, M., (Ed.) (1996b). Phonological Dyslexia [Special issue]. *Cognitive Neuropsychology*, 13, 749–934.
- Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, 100, 589–608.
- Coltheart, M., Patterson, K., & Marshall, J. C. (Eds.). (1980). *Deep dyslexia*. London: Routledge & Kegan Paul.
- Coltheart, M., & Rastle, K. (1994). Serial processing in reading aloud: Evidence for dual-route models of reading. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1197–1211.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–256.
- Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error data. *Journal of Verbal Learning and Verbal Behaviour*, 20, 611–629.
- Farrar, W. T., Orden, G. C. V., & Hamouz, V. (2001). When SOFA primes TOUCH: Interdependence of spelling, sound, and meaning in “semantically mediated” phonological priming. *Memory & Cognition*, 29, 530–539.
- Forster, K. I., & Chambers, S. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behaviour*, 12, 627–635.
- Frederiksen, J. R., & Kroll, J. F. (1976). Spelling and sound: Approaches to the internal lexicon. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 361–379.
- Friedman, R. B. (1996). Recovery from deep alexia to phonological alexia. *Brain and Language*, 52, 114–128.
- Frost, R. (1998). Toward a strong phonological theory of visual word recognition: True issues and false trails. *Psychological Bulletin*, 123, 71–99.
- Garrett, M. F. (1976). Syntactic processes in sentence production. In R. J. Wales & E. Walker (Eds.), *New approaches to language mechanisms*. Amsterdam: North-Holland.
- Glosser, G., & Friedman, R. B. (1990). The continuum of deep/phonological alexia. *Cortex*, 26, 343–359.
- Graham, K. S., Hodges, J. R., & Patterson, K. (1994). The relationship between comprehension and oral reading in progressive fluent aphasia. *Neuropsychologia*, 32, 299–316.
- Harm, M., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from a connectionist model. *Psychological Review*, 106, 491–528.
- Harm, M. W. (1998). Division of labor in a computational model of visual word recognition. Unpublished doctoral dissertation, Department of Computer Science, University of Southern California, Los Angeles, CA.
- Hendriks, A. W., & Kolk, H. H. J. (1997). Strategic control in developmental dyslexia. *Cognitive Neuropsychology*, 14, 321–366.
- Herdman, C. M. (1992). Attentional resource demands of visual word recognition in naming and lexical decisions. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 460–470.
- Hillis, A. E., & Caramazza, A. (1991). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, 114, 2081–2094.
- Jared, D. (1997). Evidence that strategy effects in word naming reflect changes in output timing rather than changes in processing route. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1424–1438.
- Jared, D., & Seidenberg, M. S. (1990). Naming multisyllabic words. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 92–105.
- Kawamoto, A. H., Kello, C., Jones, R., & Bame, K. (1998). Initial phoneme versus whole word criterion to initiate pronunciation: Evidence based on response latency and initial phoneme duration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 862–885.
- Kawamoto, A. H., Kello, C. T., Higareda, I., & Vu, J. (1999). Parallel processing and initial phoneme criterion in naming words: Evidence from frequency effects on onset and rime duration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 362–381.
- Kello, C. T., & Plaut, D. C. (2000). Strategic control in word reading: Evidence from speeded responding in the tempo naming task. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 26, 719–750.
- Kello, C. T., Plaut, D. P., & MacWhinney, B. (2000). The task-dependence of staged versus cascaded processing: An empirical and computational study of stroop interference in speech production. *Journal of Experimental Psychology: General*, 129, 340–360.

- Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Lambon Ralph, M., Ellis, A. W., & Franklin, S. (1995). Semantic loss without surface dyslexia. *Neurocase*, *1*, 363–369.
- Lupker, S. J., Brown, P., & Colombo, L. (1997). Strategic control in a naming task: Changing routes or changing deadlines? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 570–590.
- Manis, F. R., Seidenberg, M. S., Doi, L. M., McBride-Chang, C., & Peterson, A. (1996). On the bases of two subtypes of developmental dyslexia. *Cognition*, *58*, 157–196.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.
- Meyer, D. E., Osman, A. M., Irwin, D. E., & Kounios, J. (1988). The dynamics of cognition and action: Mental processes inferred from speed-accuracy decomposition. *Psychological Review*, *95*, 183–237.
- Monsell, S., Patterson, K., Graham, A., Hughes, C. H., & Milroy, R. (1992). Lexical and sublexical translation of spelling to sound: Strategic anticipation of lexical status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 452–467.
- Morais, J., Bertelson, P., Cary, L., & Alegria, J. (1986). Literacy training and speech segmentation. *Cognition*, *24*, 45–64.
- Morais, J., Cary, L., Alegria, J., & Bertelson, P. (1979). Does awareness of speech as a sequence of phones arise spontaneously? *Cognition*, *7*, 323–331.
- Ollman, R. T., & Billington, M. J. (1972). The deadline model for simple reaction times. *Cognitive Psychology*, *3*, 311–336.
- Paap, K. R., & Noel, R. W. (1991). Dual route models of print to sound: Still a good horse race. *Psychological Research*, *53*, 13–24.
- Pashler, H. (2001). Perception and production of brief durations: Beat-based versus interval-based timing. *Journal of Experimental Psychology: Human Perception and Performance*, *27*, 485–493.
- Patterson, K., Coltheart, M., & Marshall, J. C. (Eds.). (1985). *Surface dyslexia*. Hillsdale, NJ: Erlbaum.
- Patterson, K., Graham, N., & Hodges, J. R. (1994). Reading in Alzheimer's type dementia: A preserved ability? *Neuropsychology*, *8*, 395–412.
- Patterson, K., & Hodges, J. R. (1992). Deterioration of word meaning: Implications for reading. *Neuropsychologia*, *30*, 1025–1040.
- Patterson, K., & Marcel, A. J. (1992). Phonological ALEXIA or PHONOLOGICAL alexia?. In J. Alegria, D. Holender, J. Junça de Morais, & M. Radeau (Eds.), *Analytic approaches to human cognition* (pp. 259–274). New York: Elsevier.
- Patterson, K., Suzuki, T., & Wydel, T. N. (1996). Interpreting a case of Japanese phonological alexia: The key is in phonology. *Cognitive Neuropsychology*, *13*, 803–822.
- Perfetti, C. A., & Sandak, R. (2000). Reading optimally builds on spoken language: Implications for deaf readers. *Journal of Deaf Studies and Deaf Education*, *5*, 32–50.
- Plaut, D. C. (1995a). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, *17*, 291–321.
- Plaut, D. C. (1995b). Semantic and associative priming in a distributed attractor network. In *Proceedings of the 17th annual conference of the cognitive science society* (pp. 37–42). Hillsdale, NJ: Erlbaum.
- Plaut, D. C. (1997). Structure and function in the lexical system: Insights from distributed models of naming and lexical decision. *Language and Cognitive Processes*, *12*, 767–808.
- Plaut, D. C., & Booth, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychology Review*, *107*, 786–823.
- Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*, *15*, 445–485.
- Plaut, D. C., & Kello, C. T. (1999). The interplay of speech comprehension and production in phonological development: A forward modeling approach. In B. MacWhinney (Ed.), *The emergence of language* (pp. 381–415). Mahwah, NJ: Erlbaum.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.
- Rastle, K., & Coltheart, M. (1999). Serial and strategic effects in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 482–503.
- Rumelhart, D. E., Durbin, R., Golden, R., & Chauvin, Y. (1995). Backpropagation: The basic theory. In Y. Chauvin & D. E. Rumelhart (Eds.), *Back-propagation: Theory, architectures, and applications* (pp. 1–34). Hillsdale, NJ: Erlbaum.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.
- Stone, G. O., Vanhoy, M., & Van Orden, G. C. (1997). Perception is a two-way street: Feedforward and feedback phonology in visual word recognition. *Journal of Memory and Language*, *36*, 337–359.
- Stone, G. O., & Van Orden, G. C. (1993). Strategic control in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 744–774.
- Strain, E., Patterson, K., Graham, N., & Hodges, J. R. (1998). Word reading in Alzheimer's disease: Cross-sectional and longitudinal analyses of response time and accuracy data. *Neuropsychologia*, *36*, 155–171.
- Tabossi, P., & Laghi, L. (1992). Semantic priming in the pronunciation of words in two writing systems: Italian and English. *Memory & Cognition*, *20*, 303–313.
- Taylor, T. E., & Lupker, S. J. (2001). Sequential effects in naming: A time-criterion account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 117–138.
- Van Orden, G. C., & Goldinger, S. D. (1994). Interdependence of form and function in cognitive systems explains perception of printed words. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1269.

- Van Orden, G. C., Haar, M. A., Jansen op de, & Bosman, A. M. T. (1997). Complex dynamic systems also predict dissociations, but they do not reduce to autonomous components. *Cognitive Neuropsychology*, *14*, 131–165.
- 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.
- Williams, R. J., & Peng, J. (1990). An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, *2*(4), 490–501.
- Zevin, J. D., & Balota, D. A. (2000). Priming and attentional control of lexical and sublexical pathways during naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 121–135.
- Zorzi, M. (1999). Routes, races, and attentional demands in reading: Insights from computational models. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the twenty first annual conference of the cognitive science society*. London: Springer-Verlag.
- Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud. A connectionist “dual-process” model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1131–1161.
- Zorzi, M., Perry, C., Ziegler, J., & Coltheart, M. (1999). Category-specific deficits in a self-organizing model of the lexical-semantic system. In D. Heinke, G. W. Humphreys, & A. Olson (Eds.), *Connectionist models of cognitive neuroscience* (pp. 137–148). London: Springer-Verlag.