

Supporting the Serial in the SERIOL Model

Carol Whitney

Unaffiliated

629 Piping Rock Dr.

Silver Spring, MD 20905

301-384-3391

cwhitney@cs.umd.edu

Short Title: Serial Lexical Access

Abstract

It is commonly assumed that orthographical lexical access in visual word recognition takes place in parallel, with all letters activated at the same time. In contrast, in the SERIOL model of letter-position encoding, letters fire sequentially (Whitney, 2001). I present further support for such seriality on several fronts. (1) The reasons that led to the rejection of serial encoding are shown to be invalid, and the virtues of a serial encoding are discussed. (2) The SERIOL model's serial mechanisms provide a natural account of counterintuitive letter-perceptibility patterns, and correctly predict the temporal evolution of these patterns. (3) Via simulations, I show that serial lexical access accounts for conflicting data which indicate both a presence and absence of positional effects at the lexical level. In contrast, the experimental results at the letter and word levels are difficult to explain under the assumption of parallel lexical access.

Acknowledgments

I thank Jonathan Grainger for sharing his data, and for stimulating conversations. I also thank Gordon Legge, Stephen Mansfield, and Susana Chung for sharing their data, with appreciation to Stephen Mansfield for preparing that data.

Introduction

In recent years, there has been a growing interest in the issue of how letter order is encoded during visual word recognition. The SERIOL model of orthographic processing (Whitney & Berndt, 1999; Whitney, 2001, 2004) claims that letter position is represented serially. This proposal is highly controversial, as it is generally accepted that letter strings are processed in parallel.

The goal of this article to defend the claim of serial processing. More specifically, I offer such support on three fronts. (1) I discuss serial encoding of letter order in general. I examine arguments against a seriality and show that they are not valid. The benefits of a serial encoding are presented, and evidence for seriality is reviewed. (2) I then argue for SERIOL mechanisms in particular. I show that counterintuitive perceptual patterns at the letter level are explained by SERIOL's proposal for how the serial encoding is induced. (3) I then consider priming results at the lexical level. A simulation is presented that demonstrates both weak and strong positional effects, in line with experimental data.

In the following, general arguments for a serial encoding are presented first. Next, the SERIOL model is reviewed. Then the second and third topics above are presented in turn; I review the relevant experimental results, and show how the serial mechanisms of the SERIOL model account for the data.

Discussion of Serial Encoding

First I clarify what is meant by serial encoding. The heart of the matter lies in the activation of letter representations. Are all letters activated at the same time, or are the letters activated in sequence? The latter possibility has two variants. Letters may be activated strictly serially, with only one letter active at a time. The SERIOL model advocates such seriality. Alternatively, a letter may continue firing even after subsequent letters start to fire. By the time that the final letter fires, all letters are firing. The SOLAR model advocates this type of sequentiality; letter order is encoded by activation level, not

directly by order of firing (Davis, 1999). In the following, I will specifically consider strict seriality. However, many of the following arguments are also applicable sequential activation of letters in general.

Another issue involves processing routes. The encoding of letter order not only supports a direct orthographic route to lexical information, it also subserves a sublexical, phonological route. While there is some controversy concerning the capabilities of such a sublexical route (Coltheart et al., 2001; Harm & Seidenberg, 1999), models of visual word recognition have converged on the assumption that two routes are necessary in order to explain patterns of developmental and acquired dyslexia. Most models assume that a parallel encoding of letter order subserves both routes. However, Coltheart's dual-route model assumes a parallel encoding for the lexical route, and a serial encoding for the sublexical route. In contrast, I propose a serial encoding for both the lexical and sublexical routes. In the following, I will concentrate on the lexical route, and use the term *lexical access* to refer to lexical activation via that route. If referring to the sublexical route, I will denote it explicitly.

Early accounts of string processing widely assumed that the visual image was read out serially (Gough, 1972; Harcum & Nice, 1975; Mewhort, Merikle & Bryden, 1969; Sperling, 1963). However, today it is generally assumed that all letters are activated in parallel and that lexical access occurs in parallel (McClelland & Rumelhart, 1981; Harm & Seidenberg, 1999; Coltheart, Rastle, Perry, Langdon & Ziegler, 2001; Perry, Ziegler & Zorzi, 2007). First, I discuss why the shift toward the assumption of parallel processing occurred, and demonstrate that these reasons are not well motivated.

Why not Serial?

It seems that a combination of factors has contributed to the abandonment of the assumption of serial processing along the direct route to the lexicon.

Absence of Length Effect

It is generally assumed that the question of serial versus parallel processing can be adjudicated on the basis of reaction times (RTs). If there is no effect

of the number of items to be processed, this is taken as evidence of parallel processing. If RTs increase with the number of items, this is taken to reflect serial processing. When lexical-decision experiments were performed in which frequency was well controlled, no effect of string length emerged (Frederiksen & Kroll, 1976; Richardson, 1976). Therefore, these results have been taken as evidence for parallel lexical access.

However, this assumption is not necessarily warranted. As pointed out by Whitney and Lavidor (2004), serial processing could fail to yield a length effect if increased length also has a counterbalancing facilitative effect. That is, an increased number of letters may reduce the time required for the lexical network to settle after the final letter fires (as compared to shorter words). Thus increased letter-processing time and decreased lexical-settling time may cancel each other out, giving no length effect despite a serial encoding.

In fact, a recent EEG study of length effects in lexical decision provides support for this scenario (Hauk & Pulvermuller, 2004). Word length had no effect on RTs, but yielded complementary effects on EEG amplitude at different time periods. From 100 to 125 ms post-stimulus, longer words gave increased activity; from 150 to 360 ms, longer words yielded *decreased* activity. These results are entirely consistent with the proposal that longer words induce more efficient processing at the lexical stage, cancelling out increased processing time at the letter stage.

This account is also consistent with a recent analysis of behavioral length effects based on the English Lexicon Project, a on-line database of lexical-decision RTs for over 40,000 words (Balota, Cortese, Sergent-Marshall & Spieler, 2004; New, Ferrand, Pallier & Brysbaert, 2006). Once the effects of frequency, number of syllables and orthographic-neighborhood size were factored out, RTs actually *decreased* with increasing length for words of three to five letters (and remained constant for words of five to eight letters). This finding supports the idea that an increasing number of letters could have facilitative effects. For words of three to five letters, it appears that the facilitative influence outweighs the inhibitory effect of processing more letters. This pattern of length effects is replicated in the simulations below, demonstrating the plausibility of this account.

Furthermore, we have demonstrated that a length effect can be abolished (Whitney & Lavidor, 2004). It is well known that for unilateral presentation, there is a length effect in the LVF, but not the RVF (Bouma, 1973; Ellis, Young & Anderson, 1988; Young & Ellis, 1985). This difference is present even if the location of the first letter is held constant as string length is increased, indicating that it is not related to the acuity of the initial letter (Young & Ellis, 1985). Therefore, this asymmetry has been taken to reflect different modes of lexical access in the two hemispheres, with inefficient, serial RH processing and efficient, parallel LH processing (Ellis et al., 1988; Young & Ellis, 1985). However, based on predictions from the SERIOL model, we have succeeded in abolishing the LVF/RH length effect in lexical decision by using positional manipulations of contrast level (Whitney & Lavidor, 2004).

Such an adjustment could not have converted serial processing into parallel processing. Therefore, the absence of a length effect cannot be used to infer the absence of serial processing. Rather, it is more likely that LVF presentation normally yields a non-optimal orthographic encoding, which negates the usual settling advantage for longer words and reveals the underlying seriality; the contrast manipulation creates an improved activation pattern, which restores the settling advantage, thereby abolishing the LVF length effect.

In sum, reaction times are not solely a function of how long it takes to activate letter representations. These experimental results indicate that serial lexical access is not inconsistent with the absence of a behavioral length effect.

Word Superiority Effect

Using a forced choice task, Reicher (1969) demonstrated that report accuracy is better for a letter within a word than for a single letter. This phenomenon was dubbed the Word Superiority Effect (WSE). Rayner and Pollatsek (1989) argued that the WSE rules out serial lexical access, because letter-by-letter activation would imply that the non-initial letters of a word should take longer to identify than a single letter. Therefore, such letters should be at a disadvantage, rather than at an advantage.

However, the original studies on the WSE used a mask of the same size for the word and letter stimuli. Jordan and de Bruijn (1993) showed that when mask width is proportional to stimulus width, the WSE reverses; isolated letters are identified more accurately than letters within words. Furthermore, studies of word and consonant string perception at brief exposures (< 100 ms) have shown decreasing letter identification accuracy with increasing string position (Humphreys, Evett & Quinlan, 1990; Montant, Nazir & Poncet, 1998; Gomez, Ratcliff & Perea, 2007; Tydgate & Grainger, 2007). Interestingly, no advantage for the final letter is observed at these exposure durations, as discussed in more detail below. Thus the perceptual data actually do show a disadvantage for non-initial string positions, with the largest deficit for the final position. This is consistent with serial, not parallel, processing.

Rate-Coding Models

The bias toward parallel processing was also strengthened by the influence of the Parallel Distributed Processing (PDP) books (Rumelhart & McClelland, 1986). These volumes employed models that were based on the rate-coding hypothesis, wherein information is coded via average firing rate, rather than the precise timing of individual spikes.

In general, the rate-coding hypothesis was motivated by neural-recording experiments which showed that firing patterns were not reproducible across trials (Adrian, 1928). Thus, it was thought that the mechanisms underlying spike timing were inherently noisy, and therefore spike timing could not be used to encode information (Burns, 1968). Therefore, it was thought that information was encoded in the average firing rate.

The PDP volumes had an enormous impact, reviving interest in computational modeling by showing how the back-propagation algorithm could be used to train multi-layer networks. The assumptions of rate coding and distributed parallel processing came to dominate the field. Furthermore, the success of the Interactive Activation Model (McClelland & Rumelhart, 1981) probably contributed to a parallel-processing bias in the area of visual word

recognition, in particular.

However, recent experiments using more naturalistic stimuli have demonstrated that the assumption of unreliable spike generation is unwarranted. For temporally varying stimuli, spike timing is highly reproducible across trials, at the scale of a single millisecond (Mainen & Sejnowski, 1995; Rieke, Warland, de Ruyter von Stevenink & Bialek, 1997). Reproducible spike timing has also been shown in the representation of static visual images in short-term memory (Lee, Simpson, Logothetis & Rainer, 2005). Thus, contrary to earlier assumptions, neural spiking mechanisms are highly precise. Consistent with this finding, recent theories of perception have focused on the role of spike timing and oscillatory activity in the processing of visual information (Ward, 2003; VanRullen & Thorpe, 2002).

Thus one of the underlying motivations for the assumption of parallel processing - that information is represented via rate codes due to unreliability of neural spiking dynamics - proved unwarranted. Yet the bias toward parallel processing, driven in part by the impact of PDP research, remains.

Why Serial ?

Having refuted arguments *against* serial lexical access, I next discuss evidence and arguments *for* seriality.

Experimental evidence for serial readout

As we have seen, RTs cannot be used to determine whether a string's letters are processed serially or in parallel. However, seriality can be investigated by using time directly. Harcum and Nice (1975) used this approach in an experiment in which two six-letter pseudoword strings were sequentially presented, and the task was to report the letters perceived. The first string was presented for 70 ms, and the second string was presented for 35 ms. Temporal order significantly interacted with letter position. Subjects tended to report the first letter of the first string, and the second through sixth letters of the second string. Such an interaction was present for nine of the 10 subjects, while the position of the crossover point between the two strings varied with

subject.

As these strings were not words, it is unlikely that these results arose from interactions between the two strings at the lexical level. However, even if there were top-down lexical influences, this would not explain the positional interaction with presentation order. Rather, these results provide strong evidence for serial readout of the visual image(s). There was only enough time to process the initial letter(s) of the first string before the stimulus changed to the second string.

Binding

Next I discuss representational benefits of a serial encoding. We first consider the issue of how positional information could be bound abstract letter representations.

In a series of experiments using consonant trigrams, Peressotti and Grainger (1995) demonstrated that priming occurs across string positions and retinal locations. These results provide evidence for the existence of *position-independent* letter detectors, which are activated by any occurrence of a given letter, independent of its retinal location or string position. This conclusion is further supported by the existence of letter-position dyslexia in Hebrew subjects suffering from left occipitoparietal lesions. These subjects' errors were predominately anagrams of the target word (Friedmann & Gvion, 2001), indicating that the encoding of letter identity can be separated from the encoding of position.

Position-independent letter detectors require a mechanism for the dynamic binding of positional information. Note that serial firing fulfills this requirement; positional information is encoded by the order of firing. Alternatively, position could be represented by a parallel activation gradient, as in the SOLAR model (Davis, 1999). That is, activation decreases across the string, and position is encoded by activation level.

Next we consider which encoding is more compatible with experimental evidence. Under an activation gradient, the final letter would have the lowest activation level. However, for exposure durations of greater than 100 ms,

performance generally decreases with string position, but increases for the final letter (Lefton, Fisher & Kuhn, 1978; Mason, 1982); that is, the well-known final-letter advantage is observed. An activation gradient, in which the final letter is the least activated, seems inconsistent with a final-letter advantage, but this issue needs to be examined in more detail.

Perhaps performance is unrelated to degree of activation at the letter level. Rather, performance may be determined by robustness of activity at a sub-letter featural level. Indeed, a common explanation for the initial- and final-letter advantages is a general reduction in low-level lateral inhibition. However, this explanation cannot be correct, because centrally presented strings of non-alphanumeric symbols do not show an advantage for the exterior symbols; the initial and final symbols are the least well perceived, in line with the fact that they have the lowest acuity (Hammond & Green, 1982; Mason, 1982; Tydgate & Grainger, 2007). Therefore, an explanation of positional letter patterns cannot be based on general principles of low-level representations, because non-alphanumeric characters give radically different results than letters.

Rather, an explanation of perceptual patterns based on a letter-level activation gradient comes closer to explaining the data (than a lower level account), except for the final-letter advantage. It is unclear what mechanism would give a perceptual pattern that mirrors the activation gradient, except at the final letter. Davis (1999) suggests that this pattern could reflect the combination of a primacy gradient and a recency advantage, as seen in list recall. However, these phenomena are observed under strictly serial presentation; it is unclear how they would apply to SOLAR model, where there is a parallel representation once all letters have been activated.

In contrast, we will see below how the particular serial mechanism of the SERIOL model yields letter activations that generally decrease across the string, but rise for the final letter. Furthermore, this final-letter advantage depends on exposure duration, consistent with the previous observation that the final-letter advantage is not present for exposures of less than 100 ms.

Location Invariance

Next we consider the related issue of location invariance. It is well known that early visual representations are retinotopic. Activation of position-independent letter representations requires abstraction away from the encoding of retinal location. A serial encoding provides this functionality; space (i.e., a retinotopic encoding) is mapped into time (i.e., a serial encoding), resulting in a location-invariant representation.

An alternative approach to forming location-invariant representations is to assume that receptive-field size gradually increases (Mozer, 1991; Dehaene, Cohen, Sigman & Vinckier, 2005). At each successive level, locational information is reduced, while the complexity of detected features is increased. Could this approach be used instead to form a location-invariant, letter-level representation? No, if locational information were lost, there would be no information about letter position. For example, consider a detector that responds to the occurrence of a “C” across a range of retinal locations. Once this “C” detector is activated, information about the location/position of the “C” is no longer available, as there is no way to dynamically bind such information in this scenario. In sum, if you have position-independent letter detectors, there must be some way to bind locational/positional information, and any such mechanism must go beyond the proposal that location-invariance is achieved simply by increasing receptive-field size.

Phonology and Spelling

Thus far, I have focused on the direct, orthographic route to the lexicon. However, as pointed out by Goswami and Ziegler (2006), the sublexical, phonological route also requires an encoding of letter order. A serial encoding of letter order provides an abstract representation that could provide input to both reading routes. Indeed, others have presented evidence that phonology is assembled serially during visual word recognition, implying serial activation of orthographic units (Perry & Ziegler, 2002; Roberts, Rastle, Coltheart & Besner, 2003).

Also, in models of spelling (Shallice, Glasspool & Houghton, 1995; Glasspool

& Houghton, 2005), sequential input is required for training. A serial encoding of letter order directly provides the proper stimulus for this learning. Moreover, this explains why English-speaking compensated developmental dyslexics do not become good spellers, under the assumption that compensated dyslexics use a parallel slot-encoding of letter position, rather than the normal serial representation (Whitney & Cornelissen, 2005). This parallel encoding would support reading, but would not provide the proper representation for learning to spell.

Summary

The abandonment of the serial assumption seems to have been based on unwarranted interpretations of lexical-decision data (Frederiksen & Kroll, 1976; Richardson, 1976) and the WSE (Reicher, 1969), and on a general bias toward parallel processing, which was largely driven by the popularity of a certain style of modeling (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1986). Thus, there is no evidence *against* serial lexical access.

In contrast, there is direct evidence for serial read out of letter strings (Nice & Harcum, 1976). Additionally, a serial encoding provides binding and abstraction mechanisms, allowing the formation of an abstract location-invariant representation of letter order based on position-independent letter units. Such a representation also provides suitable input for the sublexical route and for spelling acquisition, and is consistent with evidence for serial phonological assembly. Thus, there are good arguments and compelling evidence *for* a serial encoding.

Review of SERIOL Model

Having discussed the motivations for a serial encoding, I now review the SERIOL model (Whitney & Berndt, 1999; Whitney, 2001, 2004) in more detail. This will allow explanations of how the serial mechanism accounts for specific patterns of experimental data.

The SERIOL framework (Sequential Encoding Regulated by Inputs to

Oscillations within Letter units) is a comprehensive account of string processing in the proficient reader. It offers a computational theory of how a retinotopic representation is converted into an abstract representation of letter order. The SERIOL model is a verbal one that quantifies activation patterns at each processing layer, and specifies how the patterns at one layer are transformed into the patterns at the subsequent layer. SERIOL has not been implemented in full, although a number of simulations based on portions of the model have been presented (Whitney & Berndt, 1999; Whitney, 2001, 2004).

The SERIOL model consists of five layers: edge, feature, letter, bigram, and word. Each layer is comprised of processing units called *nodes*, which represent groups of neurons. Within the edge, feature and letter layers, the activation of a letter is taken to be the total amount of activity occurring across all nodes representing that letter. Thus, activation depends on the number of such nodes, their firing rate, and their firing duration.

Briefly, the activation pattern at the lowest level of the model, the edge layer, corresponds to visual acuity. At the feature layer, a monotonically decreasing activation gradient is formed across the string. This activation gradient interacts with oscillatory letter nodes to produce serial firing within the letter layer. At the bigram layer, nodes recognize letter pairs that fire in a particular order. Next, I discuss this processing in more detail. For brevity, motivations for the specified representations are not reviewed. Interested readers can refer to Whitney and Berndt (1999) and Whitney (2001) for a discussion of these issues.

Edge to Feature Layers

The lowest layer of the model, the edge layer, corresponds to the early visual cortical areas. Properties of the edge level reflect the known architecture of the visual system. The edge layer is retinotopically organized, and is split along the vertical meridian, corresponding to the two cerebral hemispheres. The number of edge nodes representing a letter is taken to decrease as eccentricity increases, corresponding to the acuity gradient. Therefore, letter

activations within the edge layer decrease as eccentricity increases. For example, for the centrally presented string *cart*, A and R have a higher activation level than C and T.

The feature layer of the model is also retinotopically organized and split across the hemispheres. The model proposes that the visual system has learned to produce a monotonically decreasing activation gradient, dubbed the *locational gradient*, across this layer. In our example, C attains the highest activation, A the next highest, etc.

Note that the acuity pattern across the letters in the RVF/LH is the same as the locational gradient; both gradients *decrease* from left to right. Therefore, the acuity gradient can serve as the locational gradient in the RVF/LH. However, in the LVF/RH, the acuity gradient *increases* from left to right. Therefore, in the LVF/RH, the acuity gradient must be inverted as the edge layer activates the feature layer. Next we consider the details of this processing, which is assumed to be learned during reading acquisition in response to a top-down attentional gradient (Whitney & Cornelissen, 2005).

The locational gradient is created via modification of firing rates, as displayed in Figure 1. I propose that the visual system has learned to activate letter features more highly in the LVF/RH than the RVF/LH, via stronger edge-to-feature connection weights in the RH. This allows the first letter to reach a high level of activation. Within the RH feature layer, I propose that strong left-to-right lateral inhibition has been learned. That is, a feature node inhibits nodes to its right. As a result, features comprising the first letter receive no lateral inhibition, and inhibition increases as letter position increases. Thus, the first letter attains the highest activation level (due to strong excitation and lack of lateral inhibition), and activation decreases toward fixation (due to sharply increasing lateral inhibition as the number of letters on the left increases).

In the RVF/LH, the acuity gradient can serve as the locational gradient. Overall excitation is weaker than to the LVF/RH. Left-to-right inhibition is not necessary, although some weak such inhibition may steepen the slope of the gradient.

The two hemispheric gradients are “spliced” together via cross-hemispheric

inhibition. The RH features inhibit the LH features, bringing the activation of the LH features lower than the activation of the least activated RH features. The result is an activation gradient that strictly decreases from left to right.

Feature to Letter Layers

Next we consider how the locational gradient induces serial firing at the letter layer. Hopfield (1995) and Lisman and Idiart (1995) have proposed related mechanisms for precisely controlling spike timing, in which nodes undergo synchronous, sub-threshold oscillations of excitability. The amount of input to these nodes then determines timing of firing with respect to this oscillatory cycle. The interaction of the feature and letter layers of the model are based on this proposal.

The letter nodes comprise the oscillatory nodes, which are taken to have a cycle length of about 200 ms (i.e., in the theta range). It is assumed that the oscillatory phase is time-locked with the activation of the feature layer, such that the feature level becomes active during the trough of the oscillatory cycle. Such an assumption is consistent with MEG evidence that theta phase is reset by anticipation of stimulus presentation (Tesche & Karhu, 2000), indicating that there are brain mechanisms to coordinate input and oscillatory phase. (Under normal reading conditions, theta phase may be coordinated with saccade generation, for example.)

Input level then determines how early in the cycle a letter node can cross threshold and fire. (See Figure 2.) Near the beginning (trough) of the cycle, excitability is low, so only a letter node receiving a high level of input can cross threshold. Excitability increases over time, allowing letter nodes receiving less and less input to progressively fire. In our CART example, the C node would receive the most input, A the next, R the next, and T the least, allowing C to fire the earliest, A next, R next, and finally T.

An activated letter node inhibits other letter nodes. As a letter node continues to fire, its firing rate slows, reducing lateral inhibition to the other nodes. This allows a new letter node to start firing. When an active letter

node receives lateral inhibition, it is strongly inhibited, so that it does not re-fire for the remainder of the oscillatory cycle.

This process creates varying activation levels within the letter layer. The activation level of a letter node is determined by its firing rate and duration. Under the assumptions that a higher input level leads to faster firing and that firing duration is fairly constant across letters, the result is a decreasing activation gradient at the letter level. However, the node representing the final letter is not inhibited by a subsequent letter. It can continue fire until the end (down-phase) of the oscillatory cycle. Therefore, the final letter can fire longer than the internal letters, and reach a higher level of activation even though it receives less input. This explains the final-letter advantage.

The serial encoding provides input to both the lexical and sublexical routes. The remainder of the model specifies processing specific to the lexical route.

Letter to Bigram Layers

A bigram node XY becomes activated when letter node X fires, and then letter node Y fires within a certain time period. A bigram node responds with a burst of firing. The number of spikes in this burst decreases as the time increases between the firing of X and Y. That is, the activation of bigram XY is highest when triggered by contiguous letters, and decreases as the number of intervening letters increases. Priming data indicates that the maximum separation is likely to be two letters (Schoonbaert & Grainger, 2004).

In previous articles on the SERIOL model, I have assumed that bigram activation levels were also influenced by letter activation levels, which depend on string position (Whitney, 2001; Whitney & Berndt, 1999). However, this assumption is inconsistent with new evidence on weak positional effects of priming at the word level (Grainger, Granier, Farioli, van Assche & van Heuven, 2006), as discussed below. Therefore, I now take bigram activation levels to be affected only by the separation of the constituent letters (Whitney, 2004).

Following the evidence for a special role for external letters (Humphreys et al., 1990), the string is anchored to these endpoints via edge bigrams.¹ That is, bigram $*X$ is activated when letter X is preceded by a space, and bigram $Y*$ is activated when letter Y is followed by a space. In contrast to other bigrams, an edge bigram cannot become partially activated (e.g., by the second or next-to-last letter).

Because letters are activated sequentially, bigram activations occur sequentially. For example, the input *cart* first activates bigram node $*C$ (when letter node C fires), then CA (when A fires), then AR and CR (when R fires), then RT , AT , and CT (when T fires), and then $T*$. Due to the effects of temporal separation, bigrams $C*$, CA , AR , RT , and $T*$ attain the maximal activation level, bigrams CR and RT attain a moderate activation, and bigram CT attains a low activation.

Bigram to Word Layers

Bigram nodes connect to word nodes via weighted connections. The weight on a bigram-word connection is proportional to the activation level of that bigram when that word is presented as input (as would result from Hebbian learning). As is usual in neural network models, the weight vector is normalized, so that bigrams making up shorter words have higher connection weights than bigrams making up longer words. For example, this allows the string *tee* to activate the word node TEE more than the word node $TEETHE$.²

The input to a word node is the dot product of the bigram activations and the weights. The input changes over time, because bigram activations occur serially, as indicated above. The activation of a word node at time t is a function of its activation at time $t - 1$ and the input at time t . Lateral inhibition within the word layer also operates over time. That is, as the bigrams fire, there is ongoing competition between word nodes.

¹This is a new assumption. The importance of the external letters was formerly captured via high activations of bigrams containing those letters. However, now that bigram activation levels do not reflect letter activation levels, edge bigrams are now assumed.

²Normalization is another new assumption. Information concerning the length of the string was formerly carried on the activation levels of bigrams containing the final letter.

Discussion

In summary, the acuity gradient is transformed into the locational gradient, via hemisphere-specific processing. The locational gradient interacts with oscillatory letter nodes to induce serial firing of letter nodes. A bigram node recognizes pairs of letter nodes that fire in a particular order. Bigram nodes contact the word level via weighted connections. As word nodes are progressively activated by the bigram nodes, there is ongoing lateral inhibition between word nodes.

While the model may seem somewhat complex, keep in mind that the goal of the model is to specify how the early retinotopic representation of a string is progressively transformed into an abstract encoding of letter order, and to do so in way that is neurobiologically plausible and is consistent with experimental data. The central serial mechanism is based on computational models proposed by leading neurobiologists (Hopfield, 1995; Lisman & Idiart, 1995), and is consistent with a wide range of evidence for the importance of oscillatory activity in the brain. It should not be surprising that a comprehensive, realistic model of brain function entails some complexity.

How does the SERIOL model fit in with other models of visual word recognition? These models have focused on different aspects of the process. The Interactive Activation model (McClelland & Rumelhart, 1981) simulated the interaction between bottom-up and top-down processing on the lexical route. The dual-route (Coltheart et al., 2001), triangle (Seidenberg et al., 1994), and CDP++ (Perry et al., 2007) models have focused on the capabilities of the lexical and sublexical routes and their interaction. The triangle and CDP+ models have addressed learning on the sublexical route. For the lexical route, all of these models use a position-specific slot-encoding, i.e., separate representations for each letter in each string position. It is now becoming recognized that this encoding is unrealistic (Grainger & Whitney, 2004). The dual-route model incorporates serial processing on the sublexical route, but does not specify how serialization occurs.

In contrast, the SERIOL model focuses specifically on the orthographic aspects of visual word recognition. It addresses processing at lower visual

levels than other models, replaces the slot encoding with a serial representation that provides input to both routes, and incorporates a non-contiguous bigram representation on the lexical route. The model does not specify top-down connections, but does not mean to rule them out. The model does not address processing along the sublexical route, other than assuming that the serial encoding is parsed into a graphosyllabic representation, which is mapped onto a phonological encoding (Whitney & Cornelissen, 2005). As for learning, simulations have addressed how formation of the locational gradient could be acquired (Whitney, 2004). Other basic representational mechanisms of the model are taken to be innate.

Other models that have focused on orthographic processing include the SOLAR model (Davis, 1999), and parallel multi-letter models (Dehaene et al., 2005; Grainger et al., 2006). For a more detailed comparison of the SERIOL and SOLAR models, see Whitney (2007); for a discussion of SERIOL versus parallel multi-letter models, see Whitney and Cornelissen (2007).

Evidence for the SERIOL Model

Having reviewed the model, I now discuss some experimental support for the particulars of the model. First I look at perceptual patterns for centrally presented strings, and then look at unilateral patterns. The following explanations go beyond the general assumption of seriality; they depend on proposals specific to the SERIOL model for *how* the serial encoding is created. In this section, descriptions of how the model accounts for the data are based on reasoning about the SERIOL theory, rather than on implementations (although relevant implementations are noted in some cases).

Central Perceptual Patterns

As indicated above, the presence of a final-letter advantage seems to depend on exposure duration. For very brief exposures (< 100 ms), a final-letter advantage is usually absent (Humphreys et al., 1990; Montant et al., 1998; Gomez et al., 2007; Tydgate & Grainger, 2007), while the final-letter

advantage is present at longer exposures (Hammond & Green, 1982; Mason, 1982).

The serial mechanism of the SERIOL model explains this dependence on exposure duration. As discussed above, the final letter usually gains an advantage at the letter layer because it is not inhibited by a subsequent letter, and can therefore fire for an extended period of time. However, very brief exposures would likely decrease activation levels at the edge and feature layers. This would decrease input levels to letter the layer, and push the firing of all letters later into the oscillatory cycle. The final letter would then start firing near the end of the cycle, and its firing would quickly be terminated. Thus, the final letter would not accrue an advantage, because it cannot fire for an extended period. Whitney and Berndt (1999) presented a simulation of the feature and letter layers that demonstrated this dependence of activation patterns at the letter layer on input levels from the feature layer.

In fact, when I originally formulated the SERIOL model, the serial oscillatory mechanism was motivated by the observation that a final-letter advantage was absent in a few studies that used very brief exposures (Humphreys et al., 1990; Montant et al., 1998). It is reassuring that the absence of a final-letter advantage has been reconfirmed in more recent experiments (Gomez et al., 2007; Tydgat & Grainger, 2007).

Unilateral Perceptual Patterns

We next consider positional perceptual patterns for consonant strings presented to a single visual hemifield. Note that the model was not explicitly designed to explain such unilateral patterns. The mechanisms in the model were chosen to solve the representational problems of converting a retinotopic representation of a fixated string into an encoding of relative letter position. These mechanisms also explain the unilateral patterns. While the following explanations may seem to require a lot of machinery, keep in mind that these mechanisms are specified by a pre-existing model.

For unilateral presentation of short strings (3 or 4 letters) at large eccentricities, a counterintuitive pattern arises. In the LVF/RH, the first letter is

the best perceived of all the letters; in the RVF/LH, the last letter is the best perceived (Bouma, 1973; Estes, Allemeyer & Reder, 1976; Legge, Mansfield & Chung, 2001). Thus, in each visual field, the letter *farthest* from fixation (where acuity is the lowest) is the most likely to be correctly reported. For example, see Figure 3.

I propose that these perceptual patterns are a direct result of visual processing that is unique to letter (and number) strings. As discussed above, the induction of the serial encoding leads to differing activation levels across the letter layer. In the following, I assume that perceptibility indexes a letter’s activation level. This is a simplification, as it essentially assumes that pre-conscious recognition of a letter always proceeds correctly, and activation level then determines availability for report. For example, it assumes that when an ‘u’ is presented, the letter node ‘U’ is correctly activated, and any inability to report ‘U’ arises from a low activation level of that letter node. Actually, activation level should be weighted by the probability of correct pre-conscious recognition, which would depend on acuity. However, I assume that for nearby letters, differences in activation levels have a larger effect than differences in acuity. This simplifying assumption is consistent with the fact that letter-perceptibility patterns radically differ from what would be expected on the basis of acuity (Hammond & Green, 1982; Mason, 1982; Tydgate & Grainger, 2007), in contrast to non-alphanumeric symbols.

3

Recall that a letter’s activation level is function of its firing rate and firing duration, which are determined by bottom-up input levels to that letter node and to the next letter node, respectively. These input levels arise from the locational gradient. As discussed next, the shape of the locational gradient varies with presentation location. These differing shapes of the locational gradient and the ensuing letter-layer firing patterns then explain the unilat-

³However, recognition probability may well have a sizable effect near fixation. As discussed in more detail in the following section, the difference in acuity between contiguous letters is largest when one of the letters falls directly at fixation. In this case, the effect of a higher recognition probability for the fixated letter may dominate. This explains why a fixated letter is perceived better than the letter just to the left (Mason, 1982; Wolford & Hollingsworth, 1974), in opposition to the usual left-to-right decrease in performance.

eral patterns.

Locational Gradient Shapes

Recall that locational-gradient formation requires differential processing across the hemispheres. In the RVF/LH, the acuity gradient serves as the locational gradient. In the LVF/RH, the acuity gradient is inverted via strong bottom-up excitation and left-to-right lateral inhibition. Because the locational gradient is formed by different mechanisms in each hemisphere, the shape of the resulting gradients vary across VFs, especially when large eccentricities are considered.

Acuity is known to decrease less quickly as distance from fixation increases (Westheimer, 1987). That is, the slope of the acuity gradient is steepest near fixation, and becomes shallower as eccentricity increases.⁴ Because the RVF/LH locational gradient is based on the acuity gradient, the RVF/LH locational gradient becomes more shallow as eccentricity increases. (See right half of Figure 4.)

In the LVF/RH, formation of the locational gradient depends on left-to-right lateral inhibition. This processing is optimized to create the locational gradient for a small number of letters near fixation. For longer strings at large eccentricities, inhibition is strong at early string positions, but becomes weak at later string positions, due to distance limits on lateral inhibition. (See left half of Figure 4).

Thus the shape of the locational gradient differs across VFs. In the LVF/RH, it is shallow; in the RVF/LH, it is initially steep and then flattens out. Indeed, this account explains observed VF differences in perceptual patterns (Wolford & Hollingsworth, 1974), as demonstrated by an implemented mathematical model (Whitney, 2001). Moreover, this theory of locational gradient shape has lead to experimentally verified predictions as to how to reverse VF asymmetries associated with word length and orthographic-

⁴This pattern is commonly misrepresented as “acuity falls off rapidly outside the fovea”, implying that acuity is uniformly high across the fovea and then falls off. This is not the case. Rather, acuity falls off most rapidly within the fovea, so that acuity is substantially reduced by the fovea / parafovea boundary.

neighborhood size (Whitney & Lavidor, 2004, 2005).

Note that the proposed stronger bottom-up excitation for the LVF/RH is not inconsistent with the well-known RVF advantage observed for lexical tasks. The LVF/RH left-to-right inhibition outweighs this excitation for non-initial positions, so that input levels are lower for non-initial LVF than RVF letters. Furthermore, the resulting LVF/RH gradient is not smoothly decreasing, providing a non-optimal encoding of letter order (especially for longer words), which yields decreased activation of the target word at the lexical level.

Account of Perceptual Patterns

For a centrally fixated string, the initial-letter advantage and final-letter advantage arise for different reasons under the SERIOL model. The initial letter has an advantage because it is not inhibited from the left in the feature layer. The final letter has an advantage because it is not inhibited from the right in the letter layer. Thus, like others, I also attribute the advantage for the external letters to a lack of lateral inhibition. However, this reduced inhibition does not arise from generally decreased masking at a very low level (as is usually assumed). Rather it arises from mechanisms specific to the conversion of a spatial array of letters into a serial encoding of order.

As we see next, these proposals explain the counterintuitive perceptibility patterns for lateralized presentation of short strings. In particular, we will consider the results of Estes and colleagues (1976), given in Figure 3. In the following, retinal location will be specified in letter-widths from fixation, with negative and positive quantities denoting the LVF and RVF, respectively. In this experiment, a mask of filler characters (all \$'s or all #'s) was initially displayed across retinal locations -9 to 9 (denoted -9::9), and then four of those characters were changed to letters. In the *no-change* condition, the remaining filler characters were not affected. In the *change* condition, the remaining filler characters changed identity (to other type of filler character) when the letters appeared. The results given in Figure 3 are for the no-change condition. We first consider this condition, and then look at the

change condition below. In the following, *primacy* will signify that a letter is perceived better than all other letters, whereas *advantage* will mean that an external letter is perceived better than the internal letters.

In the no-change condition, visual attention would be automatically focused on the locations where changes occurred, minimizing the effect of the surrounding static mask characters. Therefore, in this analysis, I assume that there is little effect of the mask characters at the feature level. Figure 5 illustrates the proposed locational gradients and letter-layer firing patterns for the larger eccentricities (-8::-5 and 5::8) under the no-change condition.

First we consider LVF presentation. Within the feature layer, the initial letter is strongly excited, and strongly inhibits letters to the left, causing them to have low activation levels. At the letter level, the first letter can fire quickly and for a relatively long time, because the next letter node receives much less input (due to the strong inhibition from the first letter at the feature layer). So the initial letter attains a high activation level. Due to the low input level to the final-letter node, it starts firing late in the oscillatory cycle. Therefore, it cannot fire longer than the other letters, and no final-letter advantage emerges. This explains the initial-letter primacy and the sharp drop-off across non-initial positions observed for -8::-5, as shown in Figure 5.

For RVF presentation, overall bottom-up excitation to the feature layer is weaker. Therefore, activation the initial letter's features is relatively low. Furthermore, there is little left-to-right inhibition and the acuity/locational gradient is quite shallow, so the activation of the second letter's features is quite close to that of the first letter. At the letter level, the firing of the first letter is rapidly cut off by the second letter, giving no initial-letter advantage. Each successive letter quickly inhibits the preceding letter, allowing the final letter to start firing early in the oscillatory cycle. Therefore the final letter can fire longer than the other letters. This explains the flat pattern across non-final positions and the final-letter primacy observed for 5::8, as shown in Figure 5.

In contrast to the larger eccentricities, the perceptual function is U-shaped for -5::-2 and 2::5. Due to higher acuity, bottom-up input is higher overall.

In the LVF/RH, this allows the final letter to start firing earlier in the cycle, creating a final-letter advantage. In the RVF/LH, the acuity/location gradient is steeper, so the difference in input to the first and second letters is larger, creating an initial-letter advantage.

The present analysis implies that it should be possible to differentially affect the initial- and final-letter primacies, because the LVF initial-letter primacy depends on lack of left-to-right lateral inhibition at the feature level, while the RVF final-letter primacy requires prolonged firing at the letter level. Next we look at experimental data that support this proposal.

Manipulating the LVF initial-letter effect

As indicated above, there were also trials in the experiment presented by Estes and colleagues (1976) in which the masking characters changed identity when the letter stimuli appeared (change condition). In this case, visual attention could not automatically be directed to the letters. Therefore, the presence of the surrounding mask characters would have a stronger effect at the feature level than in the no-change condition.

The SERIOL model predicts, however, that this effect should vary with visual field and string position. The initial-letter primacy in the LVF specifically depends on a lack of left-to-right inhibition at the feature level; increased interference at the feature level should preferentially degrade perceptibility of initial letters in the LVF. That is, external letters would be expected to suffer some general decrement in perceptibility due to increased non-specific lateral masking, but an LVF initial letter should suffer an additional disadvantage because it is no longer uninhibited by features to the left.

This is exactly what occurred. Figure 6 shows the reduction in accuracy for the change condition relative to the no-change condition. Accuracy for LVF initial positions decreased by 30 points or more, while accuracy for all other location/position combinations decreased by 15 points or less. As a result, there was no LVF initial-letter primacy under the change condition. In particular, the effect of the change condition was much stronger for an initial letter than for a final letter at -5 (LVF), but did not vary with string

position at 5 (RVF); this interaction was statistically significant (Estes et al., 1976).

Manipulating the RVF final-letter effect

The RVF final-letter primacy depends on the ability of the final letter to start firing relatively early in the oscillatory cycle, so that it can fire longer than the other letters. Therefore, if it were possible to shift the firing of all letters later into the cycle, the final letter should enjoy less of an advantage because it cannot fire for as long a time. As discussed above, an overall decrease in bottom-up activation levels would have such a shifting effect, and such a decrease could be brought about by decreasing exposure duration.

Thus, we would *not* expect to see a RVF final-letter primacy at very brief exposures, because the overall low level of input pushes the firing of the final letter late into the oscillatory cycle. As exposure duration increases, the firing of all the letters is shifted earlier and earlier into the cycle, allowing the final letter to fire longer and longer. Thus, in the RVF, a final-letter primacy should only emerge as exposure duration increases. In contrast, the activation of a *non-final* RVF letter should not change much as exposure duration increases, because the firing of a non-final letter is still quickly cut off by the subsequent letter (due the shallow acuity/location gradient).

However, in the LVF, the initial-letter primacy *should* be present at very brief exposures, because it is based on strong left-to-right inhibition at the feature level, which does not depend on temporality. As exposure duration increases, the initial letter should be the primary beneficiary because, at the feature level, the increased bottom-up input to non-initial letters is canceled by increased lateral inhibition from the initial letter.

A literature search revealed that a relevant experiment had already been performed, in which retinal location and exposure duration were systematically varied in a trigram identification task (Legge, Mansfield & Chung, 2001). However, the published data were not presented in a way that would allow evaluation of the above predictions, so I inquired about obtaining the raw data from the authors, who kindly provided it. The data were analyzed

for the two largest eccentricities (-12::-10 and -11::-9 versus 9::11 and 10::12) for brief exposures (50 ms and 80 ms) versus longer exposures (125 ms and 200 ms). This analysis did indeed reveal the predicted patterns, as shown in Figure 7.

Discussion

As discussed above, these firing patterns at the letter level directly depend on specific proposals for *how* the serial encoding is created, based on hemisphere-specific processing in the formation of the locational gradient, and the interaction of the locational gradient with the oscillatory cycle. The only premise that is required in addition to the basic principles of the model is that the acuity gradient is shallower in the parafovea than the fovea, which is a known property of the visual system. No other model of orthographic processing can explain these hemifield phenomena. Next we consider possible objections to this account, and then consider testable predictions.

For non-alphanumeric symbols, it has been shown that lateral masking is asymmetric (Chastain, 1989); for a unilateral target and a single distractor, accuracy is lower when the distractor occurs to the peripheral side of the target than when it occurs to the nasal side. Perhaps, under unilateral presentation, perceptibility is poor for the inner letters (last letter in LVF, first letter in RVF) simply because they are peripherally masked, while outer letters (first letter in LVF, last letter in RVF) are not?

This explanation is unlikely. Note that the asymmetric masking results concern a fixed target, where the location of the distractor is varied. In contrast, the inner letter of a unilateral string is considerably closer to fixation than the outer letter. As discussed above, acuity determines the perceptibility of non-alphanumeric symbols (Hammond & Green, 1982; Mason, 1982; Tydgate & Grainger, 2007); internal symbols closer to fixation are perceived better than external symbols, despite the fact that an internal symbol is masked both peripherally and nasally, while an external symbol is not masked peripherally. This finding indicates that acuity overrides any general effect of lateral masking. Therefore, it is unlikely that peripheral masking can

explain why the letters closest to fixation are poorly perceived in unilateral string stimuli.

The analysis of the change versus no-change condition from Estes et al. (1976) also demonstrates that asymmetric masking cannot explain the data. At retinal location 5, the effect of the change condition was similar for positions one and four; at retinal location -5, the effect of the change condition was much stronger for position one than four. A general asymmetric masking effect should apply similarly to both visual fields; it cannot explain the interaction of position with visual field in the effect of the change condition. Rather, this pattern is predicted and explained by the SERIOL model's proposal that there is strong left-to-right, feature-level inhibition in the LVF/RH that is specific to string processing, and that the LVF initial-letter primacy arises directly from the lack of such inhibition.

Next we consider how a letter-specific result can arise at level of processing below the letter level. Using fixated strings containing both letters and non-alphanumeric symbols, Tydgat and Grainger (2007) showed that an initial letter had an advantage while an initial symbol did not. If the initial-letter advantage originates at the feature level, how this sub-letter level "know" to create a letter-specific advantage? This could occur straightforwardly under parallel processing of the stimulus via string-specific and general visual systems. The SERIOL model assumes that the feature layer is comprised of features that are specific to letter-string processing. A stimulus would activate both letter-specific features and general features. Letter-specific features would be subject to the locational gradient, while general features would reflect acuity. Letter-specific features that then activate letter representations would show the effects of string-specific processing; in particular, there will be an advantage if the letter is the initial character of a string. If a symbol is not a letter, the letter-specific features will not activate a letter representation and there will be no letter-specific effects. Rather, the symbol will be recognized via the general visual features, where the effect of acuity predominates; an initial symbol will be at a disadvantage because its acuity is lower than the internal symbols.

How could the SERIOL account be tested? There are a number of key ex-

periments which would confirm or disprove the explanations presented here. First, it is important to demonstrate within a single experiment that the presence of a final-letter advantage under central presentation depends on exposure duration. That is, exposure duration should be systematically varied; the final-letter advantage should be absent at very brief exposures, and should emerge at longer exposures. The length of the stimulus should also be varied. The SERIOL model predicts that shorter strings should show a final-letter advantage at shorter durations than longer strings.

Second, it is important to demonstrate that perceptual patterns for non-letter symbols differ from letter strings under hemifield presentation; previous studies have only looked at perceptual patterns for symbol strings near fixation. The SERIOL account predicts that there should not be an RVF initial-character primacy and LVF final-character primacy for symbols.

Third, the proposed differing sources of the RVF initial-letter primacy and the LVF final-letter primacy should be directly tested within a single experiment, by comparing the effects of letter and non-letter flankers. Placing either type of flanker to left of an LVF initial letter should have the same effect; both flanker types should reduce perceptibility of that target letter to the same degree, due to feature-level left-to-right inhibition. In contrast, a non-letter flanker placed to the right of an RVF final letter should cause less disruption than a letter flanker, because a non-letter flanker should not curtail firing of the target letter at the letter level, unlike a letter flanker.

Bigrams Revisited

We have seen how the proposed edge-to-letter processing accounts for experimental results on letter perceptibility in non-word strings. Next I show, via simulations, how the proposed letter-to-word processing explains positional results at the lexical level.

Review of Word-Level Data

The proposal of the non-contiguous bigram representation in the SERIOL model was based on masked form-priming data (Humphreys et al., 1990; Peressotti & Grainger, 1999), which indicated that relative order is an important factor in orthographic representation. Grainger and colleagues have continued their research by systematically varying the positions of target letters included in the prime (Schoonbaert & Grainger, 2004; Grainger et al., 2006). In describing such experiments, the following notation is used. A target of length n is represented by 123... n where 1 denotes the first letter, 2 the second letter, etc. The prime is specified in terms of these numbers, with “d” representing a letter not in the target. For example, the prime “rqgdzn” for the target GARDEN is denoted 3d14d6, specifying that the first letter of the prime is the third letter of the target, the second letter of the prime is not in the target, etc.

In Grainger et al. (2006), positional effects were investigated by using targets of seven and nine letters, where primes consisted of the first four or five letters of the target (initial primes) or the last four or five letters of the target (final primes). Control primes consisted entirely of letters not in the target. Within each target length and prime length, initial and final primes provided equivalent levels of facilitation (in terms of lexical-decision RTs). Thus there was no effect of the position of the prime’s letters in the target (although a non-significant advantage of <10 ms for initial primes was consistently observed across five experiments). These results were taken as evidence of parallel activation of bigram units. Similarly, in an experiment in which an initial, middle, or final pair of letters of a seven-letter target was transposed (i.e., 2134567; 1243567 + 1235467; 1234576), all primes provided equivalent amounts of facilitation (Schoonbaert & Grainger, 2004).

However, in another experiment using primes that contained non-target letters, a positional effect did emerge (Schoonbaert & Grainger, 2004). For seven-letter targets and primes where the initial, middle, or final two letters were replaced with letters not in the target (i.e., dd34567; 12dd567+123dd67; 12345dd), only 12345dd provided facilitation. For five letter strings, there was

no effect of position, as neither dd345, 12dd5, nor 123dd produced facilitation, indicating that there was simply not enough overlap with the target (and/or too much overlap with competitors) to provide facilitation.

In the following, I will focus on priming studies of seven-letter targets, because results for targets of a single length can be straightforwardly compared, and seven letters is presumably long enough to maximize positional effects, but short enough that visual acuity is not a strongly limiting factor in performance. In keeping with the focus on seriality, the goal here is to account for the presence and absence of positional effects, not to provide a comprehensive account of all priming data.

It is of particular interest to consider the results for 12345 and 34567 versus 12345dd and dd34567. Both 12345 and 12345dd produced facilitation, while 34567 primed, but dd34567 did not. Thus it appears that distractor letters have a stronger inhibitory effect when they are near the beginning of the prime.

Another way to investigate positional effects at the lexical level is to analyze error patterns under degraded perceptual conditions. For very briefly presented word targets (< 50 ms) of six or fewer letters, letters toward the beginning of the target are more likely to be retained in erroneous responses than letters toward the end of the target (Humphreys et al., 1990; Montant et al., 1998).

A similar pattern is also present in errors made under normal presentation to aphasic patients (Whitney & Berndt, 1999), as shown in Figure 9. A letter in the n th position in the target was scored as being retained if it occurred in n th position in the response. A similar decreasing pattern is also obtained if the n th target letter can occur in any response position, but this pattern is not present if the target and response are aligned at the final letter, and scored backward. Thus there is a strong positional effect that is not an artifact of scoring method. As this pattern is similar to normals under very rapid presentation, it is likely that this pattern arises from an inherent aspect of visual word recognition, rather than from brain damage per se.

Thus initial letters have an advantage when noise is present - under aphasia, for very brief presentation to normals, and for primes containing non-

target letters. In contrast, for primes consisting only of target letters, a significant positional effect is not present.

Previous Simulation

In previous work, a simulation of the bigram and word levels (Whitney & Berndt, 1999) replicated the aphasic error pattern. Bigram activations were taken to depend on both the position and separation of the constituent letters. Aphasia was simulated by adding noise to word activations. Due to the positional dependence of bigram activations, erroneous “responses” tended to preserve the initial letters of the target.

However, this assumption incorrectly predicts that there should be a positional priming effect for initial (e.g., 12345) versus final (e.g., 34567) primes. That is, the reduced activation levels of bigrams toward the end of the word, which allowed replication of the aphasic error pattern, should also yield significantly less priming than bigrams near the start of the word. However, as discussed above, this is not the case (Grainger et al., 2006). This lack of a positional effect led to an abandonment of the assumption that bigram activation levels depend on letter activation levels, which depend on string position. Now bigram activations depend only on the separation between the constituent letters (Whitney, 2004).

Another difficulty with the original simulation is that it required an additional assumption in order to replicate aphasic error pattern. It was necessary to assume that input levels to letter nodes were reduced, so that there was no final-letter advantage. While it is plausible that aphasia would affect input levels to letter nodes, it would be more desirable to simulate the error pattern without this assumption.

Also, the original implementation did not simulate reaction times. Due to the widespread assumption that length effects can be used to determine whether processing is serial or parallel, I wished to demonstrate that serial processing could replicate the experimental findings of a reverse length effect for three to five letters and no length effect for longer words.

Simulations

Therefore, I sought to implement an improved bigram-to-word simulation that would meet four goals: (1) correct recognition of all words in a large database, with settling times reproducing experimental effects of length; (2) replication of aphasic error pattern under noise; (3) presence of positional effect for primes containing non-target letters; (4) lack of positional effect for primes containing only target letters. In the original simulation, bigrams activated the word layer in a single time step in a purely bottom-up manner; the simulation did not model the serial activation of bigrams or lateral inhibition within the lexical level. The present simulation included these temporal factors, which allowed these goals to be met without any assumptions about positional activation patterns.

Next an overview of the simulation is presented; details are given in Appendix A. The simulation consisted of two layers of nodes: the bigram layer and the word layer. The bigram layer consisted of all possible two-letter pairs, plus edge bigrams. The word level consisted of word nodes representing 3650 single-syllable words, as in the original simulation, plus 50 seven-letter words in order to simulate the priming studies.

Connections between bigrams and words were weighted, where weights were proportional to the activation of the given bigram for the given word. Weights on shorter words were larger than weights on longer words; the degree of reduction with length was controlled by a parameter.

During the simulation, bigram activations were sequentially clamped, and activation level depended on the separation of letters in the target. Separations of 0, 1, 2 and >2 letters had activations of 1.0, C_{S1} , C_{S2} , and 0, respectively, where C_{S1} and C_{S2} are parameters with $1.0 < C_{S1} \leq C_{S2}$. At each time step t , the bigrams that would be activated by the character in position t were set their corresponding values. For an example, see Table 1.

At each time step, word-level activations were then updated in two stages. (a) The incoming activation to each word node was simply added to its current activation. The incoming activation is given by the dot product of the bigram activations and the word node's weights. That is, each bigram

activation is multiplied by the corresponding weight; the incoming activation is the sum of these products. (b) Then the effect of lateral inhibition within in the word layer was simulated by re-updating each word node’s activation. The overall strength of inhibition was determined by a parameter.

Each input string was presented over $L + 1$ time steps, where L is the length of the input. For normal recognition, the simulation was continued until a settling criterion was reached (as described below). For the other simulations, the result was read off of the word nodes after the final edge bigram was presented. For aphasic recognition, the result was simply the most highly activated word node. For the priming simulations, the activation of the target word node was taken to index the degree of facilitation.

The simulations involved four parameters: two for bigram activations, one for length normalization, and one for strength of lateral inhibition. The above four goals were often at cross purposes. Goal (1) requires normalization of the weight vector. Yet if shorter words have too much of an advantage, they excessively inhibit longer words, under the inhibition required for goals (2) and (3). Goals (2) and (3) require a positional effect, while goal (4) requires no positional effect. The parameters were hand-tuned to meet these goals; the parameter settings are given in Appendix A. This single set of parameter values was used in all of the following simulations.

Normal Recognition

After the final edge bigram was presented, lateral inhibition within the word layer was continued until the most active word node attained a significantly higher activation than the other word nodes. That is, the simulation was stopped when the activation of the most active word exceeded that of the second most active word by a given amount, denoted the *differential*. Figure 8 shows the average number of time steps required under two specifications of the differential. All words in the database were correctly recognized.

In one case, the differential was set to 60% of the activation of the most active word node; here the differential decreased with word length, because maximum activation decreases with length. In the other case, the differential

was set to a constant value of 4.0. Both criteria produce a similar pattern, with settling time decreasing across lengths three to five, and flattening across lengths of five to seven. Thus, for words of three to five letters, increasing length gave a settling advantage even though more time steps were required to present the letters; this effect was present even if the differential was smaller for shorter words (i.e., in the 60% criterion). This settling advantage occurred because more letters supplied more information, allowing longer words to be more quickly differentiated from competitors. As word length increased to five to seven letters, the degree of this advantage decreased, causing the flattening of settling-time curve; this effect was present even if the differential did not increase with length (i.e., in the 4.0 criterion).

Priming

Eight prime conditions were simulated for seven-letter targets, comprising two deletion conditions (12345 and 34567), three transposition conditions (2134567, 1243567+1234657, and 1234576) and three replacement conditions (12345dd, 12dd567+123dd67, and dd34567). Recall that there was no positional effect for deletion and transposition primes experimentally, while only 12345dd provided facilitation for the replacement primes. The goal of the priming simulations was to replicate this pattern.

20 of the seven-letter words added to the database were used as targets. A prime string was generated for each target for each prime condition. Priming was simulated by presenting the prime string as input, and then reading off the activation of the corresponding target word node. For example, in a trial simulating the 12345dd condition for the target *without*, the string *withoak* was sequentially encoded over the bigram nodes. Following activation of the final edge bigram (K*), the activation of the word node WITHOUT was recorded.

To form the primes for the dd34567 and 12dd567+123dd67 conditions, letter replacements chosen so that the distractor letters did not form actual words in conjunction with nearby target letters. For the 12345dd condition, the distractor letters were chosen such that 5dd formed a word in the database

in most cases (such as *oak* in the *withoak* example). This choice of distractor letters follows the observation (Davis, pers. comm.) that distractor / target letter combinations in the experimental stimuli (Schoonbaert & Grainger, 2004) were more similar to actual words in the 12345dd than the dd34567 condition. The target words and primes for the replacement conditions are given in Appendix B.

The simulation results are given in Table 2. It is evident that non-final replacements (dd34567 and 12dd567+123dd67) yielded significantly lower scores than the other prime conditions, consistent experimental results the these conditions did not provide facilitation, while the others did.

In the simulation, non-final replacements yielded low activation levels because the target word node received early inhibition from word nodes activated by the distractor letters. Although the distractor letters did not actually form words in conjunction with the target letters, the distractor letters nevertheless caused non-target words to become partially activated, due to the flexibility of the bigram encoding. The inhibition from these competitors then put the target word at a disadvantage from which it could not recover.

Despite the fact that the distractor letters formed words in the 12345dd primes, target word-node activations were high because the target word attained an early advantage, and inhibited the later competitors. Thus the simulation results for replacement primes depend directly on the temporal sequence of bigram activations, going in the opposite direction of what would be expected on the basis of similarity to competitor words, in agreement with the experimental results.

For the deletion primes (12345 and 34567), inhibition from competitors was reduced, and the difference between initial and final primes was small. This difference is numerically consistent with the results of Grainger et al. (2005), where a non-significant advantage (< 10 ms) for initial primes was consistently observed across five experiments.

For the transposition primes (1234576, 2134567, and 1243567+1235467), there was again a relatively small difference between the initial and final manipulations. There was a larger difference between initial and middle transpositions, although experimental results showed no difference between these

conditions. However, all of the transposition primes yielded high simulated activations, which may correspond to ceiling levels of facilitation experimentally.

It is also evident in the simulated results that the non-final replacement primes yielded more variance than the other conditions. This occurred for the same reason that the simulations were able to replicate the experimental findings. In the non-final replacements, target-word activations were strongly influenced by inhibitory effects, which depended on the similarity of the stimulus to the word set. Hence, inhibitory effects varied with the stimuli. In the other conditions, there was little variance because target activation level was primarily determined by the bottom-up match between the stimulus and the word.

Aphasic Error Pattern

In evaluating the experimental data on errors made by aphasics, a letter was counted as being retained if occurred in the erroneous response in the same position as in the target word (Whitney & Berndt, 1999). More complex methods of scoring yielded the same positional patterns as this position-specific method, so the position-specific metric was used for simplicity. In order to compare the simulated data to the experimental data, I use the same metric.

A lesion was simulated by adding normally distributed noise to each word node at each time step, prior to the inhibition stage. Noise with mean 0.3 and standard deviation 0.45 yielded good results, shown in figure 9. As is evident, the probability of retaining a letter decreased with its position (in particular, across positions one through four). In contrast, backward scoring of the results gave a much flatter pattern (as was also the case for the experimental data), indicating that the forward positional effect was not merely an artifact of the scoring method. In the simulation, this positional effect arose from the sequential activation of the bigram nodes, coupled with the ongoing lateral inhibition. Potential erroneous responses that were not initially highly activated became inhibited and remained at a disadvantage.

Therefore, retention level was highest for early string positions and decreased across the string, giving a strong positional effect.

For comparison, the same simulation was also run without lateral inhibition. The results are given in Figure 10. In this case, forward and backward scoring yielded the same pattern; a weaker positional effect is evident in both cases. These positional effects are therefore an artifact of the position-specific scoring method. (As distance from an edge letter increases, the probability of position-altering insertions and deletions increases, reducing the probability of a retained letter remaining in the same position as in the target.)

The results of the lesioned simulation also showed other similarities to the experimental data. Aphasic subjects tended to preserve word length in their erroneous responses. Experimentally, average response length to targets of lengths 3, 4, 5, and 6 were 4.0, 4.2, 4.9, and 5.9, respectively (Whitney & Berndt, 1999). The simulated data also showed sensitivity to target length, giving 4.3, 4.6, 5.1, and 5.6. Retention level at a given position tended to increase with target length for both the aphasics and the simulation. For example, for position 3, experimental retention rates were 40%, 55%, 65%, and 55% for targets of lengths 3, 4, 5, and 6, respectively, while the simulation gave 21%, 43%, 68%, and 75%.

Discussion

The simulations fulfilled the stated objectives. For normal recognition, settling time replicated experimental results on the effect of word length; in particular, there was a *reverse* length effect for shorter words, even though letters were processed serially. In the other simulations, there was a strong positional effect for primes with distractor letters, a weak positional effect for primes without distractor letters, and strong positional effect for errors under noise. The differing positional patterns occurred because noise and distractor letters amplified the serial effects of lateral inhibition within the word level.

The present observation of a difference in positional effects for deletion versus replacement primes is based on results of separate experiments. In future work, it would be useful to reconfirm this finding by directly comparing

these prime types within a single experiment.

Conclusion

We have seen that there are no data that contradict a serial encoding, and have considered how a serial encoding provides important representational capacities. The pre-existing SERIOL model accounts for perceptual patterns that are otherwise difficult to explain. The unilateral patterns are explained by SERIOL's serial mechanism: the specification of how the locational gradient is formed from the acuity gradient, known properties of the acuity gradient, and the interaction of the locational gradient with the oscillatory cycle.

We have also seen that serial lexical access, in conjunction with ongoing lateral inhibition within the word level, explains positional effects in the presence of noise (i.e., distractor letters, brain damage, or very brief presentation), and the lack of positional effects in low noise. In contrast, parallel lexical access cannot explain these differing patterns.

Previously, the pendulum swung from the assumption of a serial readout of a letter string (Gough, 1972; Harcum & Nice, 1975; Mewhort, Merikle & Bryden, 1969; Sperling, 1963) to the assumption of parallel processing (Coltheart et al., 2001; McClelland & Rumelhart, 1981; Seidenberg & McClelland, 1989). Given the lack of a valid argument against a serial encoding, the direct evidence for serial processing (Nice & Harcum, 1976), the ability of a serial encoding to provide important representational functions, and the explanatory capacity of SERIOL's serial mechanism, I suggest that it is time for the pendulum to swing back to seriality, and remain there.

Appendix A

This appendix specifies the simulations. As described in the text, the simulations were carried out on a two-layer network, consisting of the (input) bigram layer and the (output) word layer. The functions implementing the length normalization and lateral inhibition were chosen on the basis of computational efficiency, rather than neurobiological plausibility. In the following, C denotes a parameter.

Let B^{xy} denote a bigram node representing the letter x followed by the letter y . A bigram's activation A for a string S is a function of the number of letters separating x and y . As specified in the text, $A(B^{xy}, S) = 1.0$ for separation = 0, C_{S1} for separation = 1, C_{S2} for separation = 2, and 0 otherwise. An input string S was presented over $L(S) + 1$ time steps, where $L(S)$ gives the length of S . At each time step t , the bigram layer was clamped to the corresponding bigram activations. (In the following, a bigram not specified has an activation of 0.) For example, at time step 1, bigram *1 was clamped to 1.0. At time step 2, bigram 12 was clamped 1.0. At time step 3, bigram 23 was clamped to 1.0 and bigram 13 to C_{S1} , etc. At final time step, bigram L^* , was clamped to 1.0.

Normalization of the weight vector is assumed, to give an advantage for shorter words. Let Wd^S represent a word node encoding string S . The weight from a bigram node to a word node is given by:

$$W(B^{xy}, Wd^S) = \frac{C_{nrm}}{L(S) + C_{nrm}} * A(B^{xy}, S)$$

If a bigram received two different activation levels for a word, the larger of $A(B^{xy}, S)$ was taken in determining the weight. This scaling of the bigram's activation level provides normalization by decreasing the weights for longer words, via division by $L(S)$. The constant C_{nrm} modulates this normalization; the higher its value, the less the effect. Thus, weights were normalized, while bigram activations were not.

Initially, all word activations were 0. At each time step, all word activations were first updated on the basis of bottom-up input; the dot-product of

the bigram activation vector and the weight vector was simply added to the current activation. Then the effects of lateral inhibition were simulated by re-updating each word's activation, A , as follows:

$$A(Wd^S, t) = A(Wd^S, t) - C_{inh} * (AMax(t) - A(Wd^S, t))^2$$

where $AMax(t)$ denotes the activation of the maximally active word node, and the constant C_{inh} scales the strength of the inhibition.

The results described in the text used the following parameter values:

$$C_{S1} = 0.8 \quad C_{S2} = 0.4 \quad C_{nrm} = 25 \quad C_{inh} = 0.08$$

Appendix B

This appendix specifies the target words and replacement primes used in the priming simulations. *Initial* refers to the dd34567 condition, *Middle* to the 12dd567+123dd67 condition, and *Final* to the 12345dd condition. All priming simulations used the same target words.

Target	Initial	Middle	Final
without	abthout	wibcout	withoak
picture	dgcture	picdlre	pictusk
subject	febject	sudfect	subjeat
program	bhogram	profham	progrip
provide	ijovide	prijide	proviar
perhaps	klrhaps	pernlps	perhaim
outside	mntside	oumnide	outsipq
nothing	pqthing	notqrng	nothire
network	rstwork	nersork	netwold
however	tuwever	howtuer	howevim
history	vustory	hivwory	histopt
further	xyrther	furwyer	furtham
stumble	ozumble	stobrlr	stumbad
contact	dfntact	condfct	contarm
control	gjntrol	comjrol	contrgj
between	hktween	bethken	betwelm
average	lnerage	avlnage	averajk
article	muticle	artmple	articop
trouble	vpouble	tranble	trouban
believe	dwlieve	belgtve	beliend

References

- Adrian, E. (1928). *The Basis of Sensation*. New York, NY: Norton.
- Balota, D., Cortese, M., Sergent-Marshall, S. & Spieler, D.H. Yap, M. (2004). The English Lexicon Project: A web-based repository of descriptive and behavioral measures for 40,481 English words and nonwords. <http://lexicon.wustl.edu>. Washington University.
- Bouma, H. (1973). Visual interference in the parafoveal recognition of initial and final letters of words. *Vision Research*, 13, 767–782.
- Burns, B. (1968). *The Uncertain Nervous System*. London, UK: Arnold.
- Chastain, G. (1989). Effects of relative position and shape relation on gap detection with geometric form pairs. *Acta Psychologica*, 72, 13–23.
- 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.
- Davis, C. (1999). *The Self-Organising Lexical Acquisition and Recognition (SOLAR) model of visual word recognition*. PhD thesis, University of South Wales.
- Dehaene, S., Cohen, L., Sigman, M. & Vinckier, F. (2005). The neural code for written words: a proposal. *Trends in Cognitive Science*, 9, 335–341.
- Ellis, A. W., Young, A. W. & Anderson, C. (1988). Modes of word recognition in the left and right cerebral hemispheres. *Brain and Language*, 35, 254–273.
- Estes, W., Allemeyer, D. & Reder, S. (1976). Serial position functions for letter identification at brief and extended exposure durations. *Perception and Psychophysics*, 19, 1–15.
- Frederiksen, J. & Kroll, J. (1976). Spelling and sound: Approaches to the internal lexicon. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 361–379.

- Friedmann, N. & Gvion, A. (2001). Letter position dyslexia. *Cognitive Neuropsychology*, 18, 673–696.
- Glasspool, D. & Houghton, G. (2005). Serial order and consonant-vowel structure in a graphemic output buffer model. *Brain and Language*, 94, 304–330.
- Gomez, P., Ratcliff, R. & Perea, M. (2007). A model of letter position coding: The overlap model. Submitted.
- Goswami, U. & Ziegler, J. C. (2006). A developmental perspective on the neural code for written words. *Trends in Cognitive Sciences*, 10, 142–143.
- Gough, P. (1972). One second of reading. In J. Kavanagh & I. Mattingly (Eds.), *Language By Ear and by Eye: The Relationships between Speech and Reading* (pp. 331–358). Cambridge: MIT Press.
- Grainger, J., Granier, J., Farioli, Q., van Assche, E. & van Heuven, W. (2006). Letter position information and printed word perception: the relative-position priming constraint. *Journal of Experimental Psychology*, 32, 865–884.
- Grainger, J. & Whitney, C. (2004). Does the human mind read words as a whole? *Trends in Cognitive Science*, 8, 58–59.
- Hammond, E. & Green, G. (1982). Detecting targets in letter and non-letter arrays. *Canadian Journal of Psychology*, 36, 67–82.
- Harcum, E. & Nice, D. (1975). Serial Processing shown by mutual masking of icons. *Perceptual and Motor Skills*, 40, 399–408.
- Harm, M. W. & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: insights from connectionist models. *Psychological Review*, 106, 491–528.
- Hauk, O. & Pulvermüller, F. (2004). Effects of word length and frequency on the human ERP. *Clinical Neurophysiology*, 115, 1090–1103.

- Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376, 33–36.
- Humphreys, G. W., Evett, L. J. & Quinlan, P. T. (1990). Orthographic processing in visual word identification. *Cognitive Psychology*, 22, 517–560.
- Jordan, T. R. & de Bruijn, O. (1993). Word superiority over isolated letters: The neglected role of flanking mask-contours. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 549–563.
- Lee, H., Simpson, G., Logothetis, N. & Rainer, G. (2005). Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron*, 45, 147–156.
- Lefton, L., Fisher, D. & Kuhn, D. (1978). Left-to-right processing of alphabetic material is independent of retinal location. *Bulletin of the Psychonomic Society*, 12, 171–174.
- Legge, G. E., Mansfield, J. S. & Chung, S. T. (2001). Psychophysics of reading. XX. Linking letter recognition to reading speed in central and peripheral vision. *Vision Research*, 41, 725–743.
- Lisman, J. E. & Idiart, M. A. (1995). Storage of 7 +/-2 short-term memories in oscillatory subcycles. *Science*, 267, 1512–1515.
- Mainen, Z. & Sejnowski, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268, 1503–1506.
- Mason, M. (1982). Recognition time for letters and non-letters: Effects of serial position, array size, and processing order. *Journal of Experimental Psychology*, 8, 724–738.
- McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375–407.

- Mewhort, D., Merikle, P. & Bryden, M. (1969). On the transfer from iconic to short-term memory. *Journal of Experimental Psychology*, 81, 89–94.
- Montant, M., Nazir, T. A. & Poncet, M. (1998). Pure alexia and the viewing position effect in printed words. *Journal of Cognitive Neuropsychology*, 15, 93–140.
- Mozer, M. (1991). *The Perception of Multiple Objects: A Connectionist Approach*. Cambridge: MIT Press.
- New, B., Ferrand, L., Pallier, C. & Brysbaert, M. (2006). Re-examining word length effects in visual word recognition: New evidence from the English Lexicon Project. *Psychonomic Bulletin and Review*, 13, 45–52.
- Nice, D. & Harcum, E. (1976). Evidence from mutual masking for serial processing of tachistoscopic letter patterns. *Perceptual and Motor Skills*, 42, 991–1003.
- Peressotti, F. & Grainger, J. (1995). Letter-position coding in random constant arrays. *Perception and Psychophysics*, 57, 875–890.
- Peressotti, F. & Grainger, J. (1999). The role of letter identity and letter position in orthographic priming. *Perception and Psychophysics*, 61, 691–706.
- Perry, C., Ziegler, J. & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: the CDP+ model of reading aloud. *Psychological Review*, 114, 273–315.
- Perry, C. & Ziegler, J. C. (2002). Cross-language computational investigation of the length effect in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 990–1001.
- Rayner, K. & Pollatsek, A. (1989). *The Psychology of Reading*. Englewood Cliffs, NJ: Prentice-Hall.
- Reicher, G. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81, 275–280.

- Richardson, J. (1976). The effects of stimulus attributes on latency of word recognition. *British Journal of Psychology*, 67, 315–325.
- Rieke, F., Warland, D., de Ruyter von Stevenink, R. & Bialek, W. (1997). *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press.
- Roberts, M. A., Rastle, K., Coltheart, M. & Besner, D. (2003). When parallel processing in visual word recognition is not enough: new evidence from naming. *Psychonomic Bulletin Review*, 10, 405–414.
- Rumelhart, D. & McClelland, J. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volumes 1 and 2*. Cambridge, MA: MIT Press.
- Schoonbaert, S. & Grainger, J. (2004). Letter position coding in printed word perception: Effects of repeated and transposed letters. *Language and Cognitive Processes*, 19, 333–367.
- Seidenberg, M. S. & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523–568.
- Seidenberg, M. S., Plaut, D. C., Petersen, A. S., McClelland, J. L. & McRae, K. (1994). Nonword pronunciation and models of word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1177–1196.
- Shallice, T., Glasspool, D. & Houghton, G. (1995). Can neuropsychological evidence inform connectionist modelling? Analyses from spelling. *Language and Cognitive Processes*, 10, 195–255.
- Sperling, G. (1963). A model for vision memory tasks. *Human Factors*, 5, 19–31.
- Tesche, C. & Karhu, J. (2000). Theta oscillations index human hippocampal activation during a working memory task. *Proceedings of the National Academy of Sciences*, 97, 919–924.

- Tydgat, I. & Grainger, J. (2007). Serial position effects in the identification of letters, digits, and symbols. In *Proceedings of XVth Meeting of the European Society for Cognitive Psychology* (p. 123).
- VanRullen, R. & Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision Research*, 42, 2593–615.
- Ward, M. (2003). Synchronous neural oscillations and cognitive processes. *Trends in Cognitive Science*, 7, 553–559.
- Westheimer, G. (1987). Visual Acuity. In R. Moses & W. Hart (Eds.), *Adler's Physiology of the Eye: Clinical Application* (pp. 531–547). St. Louis, MO: The C.V. Mosby Company.
- Whitney, C. (2001). How the brain encodes the order of letters in a printed word: The SERIOL model and selective literature review. *Psychonomic Bulletin and Review*, 8, 221–243.
- Whitney, C. (2004). *The Neural Basis of Structured Representations*. PhD thesis, University of Maryland, College Park, MD.
- Whitney, C. (2007). A comparison of the SERIOL and SOLAR theories of letter-position encoding. *Brain and Language*. In press.
- Whitney, C. & Berndt, R. S. (1999). A new model of letter string encoding: Simulating right neglect dyslexia. *Progress in Brain Research*, 121, 143–163.
- Whitney, C. & Cornelissen, P. (2005). Letter-position encoding and dyslexia. *Journal of Research in Reading*, 28, 274–301.
- Whitney, C. & Cornelissen, P. (2007). SERIOL reading. *Language and Cognitive Processes*. In press.
- Whitney, C. & Lavidor, M. (2004). Why word length only matters in the left visual field. *Neuropsychologia*, 42, 1680–1688.

- Whitney, C. & Lavidor, M. (2005). Facilitative orthographic neighborhood effects: The SERIOL model account. *Cognitive Psychology*, 51, 179–213.
- Wolford, G. & Hollingsworth, S. (1974). Retinal location and string positions important variables in visual information processing. *Perception and Psychophysics*, 16, 437–442.
- Young, A. & Ellis, A. (1985). Different methods of lexical access for words presented to the left and right visual fields. *Brain and Language*, 24, 326–358.

Table Captions

Table 1: Bigram activations at each time t for the input string *cart*. At each t , bigrams not shown have activation 0.

Table 2: Mean activation and Standard Error of the Mean for target words in priming simulations, including the identity condition for comparison.

t	1.0	C_{S1}	C_{S2}
1	*C		
2	CA		
3	AR	CR	
4	RT	AT	RT
5	T*		

Table 1:

Prime	Act.	SEM
1234567	9.61	0.01
1234576	7.86	0.09
2134567	7.35	0.05
1243567	8.41	0.05
12345	5.89	0.04
34567	5.65	0.03
12345dd	5.85	0.02
dd34567	4.19	0.22
12dd567	3.82	0.40

Table 2:

Figure Captions

Figure 1: Formation of the locational gradient at the feature layer, for the centrally fixated stimulus CASTLE. The horizontal axis represents retinal location, while the vertical axis represents activation level. The bold-face letters represent bottom-up input levels, which are higher in the RH than the LH. In each hemisphere, activation decreases as a eccentricity increases, due to the acuity gradient. The italicized letters represent the effect left-to-right inhibition within the RH, and RH-to-LH inhibition in the LH. In the RH, C inhibits A, and C and A inhibition S, creating a decreasing gradient. The RH inhibits each letter in the LH by the same amount, bringing the activation of T lower than that of S. As a result, activation monotonically decreases from left to right.

Figure 2: Interaction of input level and timing of firing for a cell with undergoing a sub-threshold oscillation of excitability. When a relatively high level of input (top curving line) is added to the base oscillation, the cell crosses threshold at time 1 (action potential not illustrated). If less input were received, the cell would cross threshold later in the cycle, such as at time 2.

Figure 3: Experimental results regraphed from Figure 3 of Estes et al. (1976) for a four-letter string, occurring across two different retinal locations (given in letter widths) in each visual field. Exposure duration was 2400 ms. (Subjects were trained to maintain central fixation, and their gaze was monitored.)

Figure 4: Schematic of locational gradients for the stimulus CART at three different presentation locations. The vertical axis represents activation, while the horizontal axis represents retinal location. For central presentation, the gradient is smoothly and rapidly decreasing. For RVF presentation, the gradient is shallower because the acuity gradient is shallower. For LVF presentation, the initial letter strongly inhibits nearby letters, but the gradient flattens out as acuity increases.

Figure 5: Locational gradients and resulting firing patterns for LVF/RH (normal font) and ***RVF/LH*** (bold italics) presentation. Top: Comparison of locational gradient for string CDFG under RVF/LH presentation and LVF/RH presentation. Bottom: Cartoon of resulting firing pattern at the letter level. The point in the oscillatory cycle at which the down phase prevents further firing is marked *. In the LVF/RH, the first letter fires faster and longer than the other letters, because it receives a much higher level of input. The variations in the amount of bottom-up input creates decreasing activation across the string. The final letter starts firing late in the cycle, and is soon cut off by the end of the oscillatory cycle, giving no final-letter advantage. In the RVF/LH, each letter rapidly cuts off the previous letter, allowing the final letter to fire a long time. As a result, activation is flat across the string and rises for the final letter. These firing patterns account for the perceptibility patterns at the larger eccentricities in Figure 3

Figure 6: Effect of the change condition relative to the no-change condition for external letters under 2400 ms exposure duration (i.e., percent correct in no-change condition minus percent correct in change condition). Values were read off of Figure 3 (Estes et al., 1976), rounding to nearest 5 percentage points. At retinal locations 2 and -2 the change condition had no effect, likely due to the high acuity.

Figure 7: Results from Experiment 2 (Legge et al., 2001) for the two largest eccentricities, grouped by exposure duration, with 95% confidence intervals.

Figure 8: Average number of time steps to settle by word length in the simulation of normal recognition. 60% denotes a differential set to 0.6 times the activation of the maximally active word node; 4.0 denotes a constant differential.

Figure 9: Experimental (Whitney & Berndt, 1999) and simulated results for the aphasic error pattern, with backward scoring of simulated results for comparison. *Percent Retained* refers to the percentage of erroneous trials in

which the letter in the i th position in the target occurred in the i th position the response ($n = 201$ for experiment; $n = 367$ for simulation). In backward scoring, the target and response were aligned at the final letter, and scored from right to left. In this case, position 1 corresponds to the final letter, 2 corresponds to the next-to-last letter, etc. Data are collapsed over target lengths of three to six. (In the both the experimental data and the simulation, there same patterns were present within each target length.)

Figure 10: Results of aphasic simulation without lateral inhibition, for forward versus backward scoring.




Figure 1:




Figure 2:

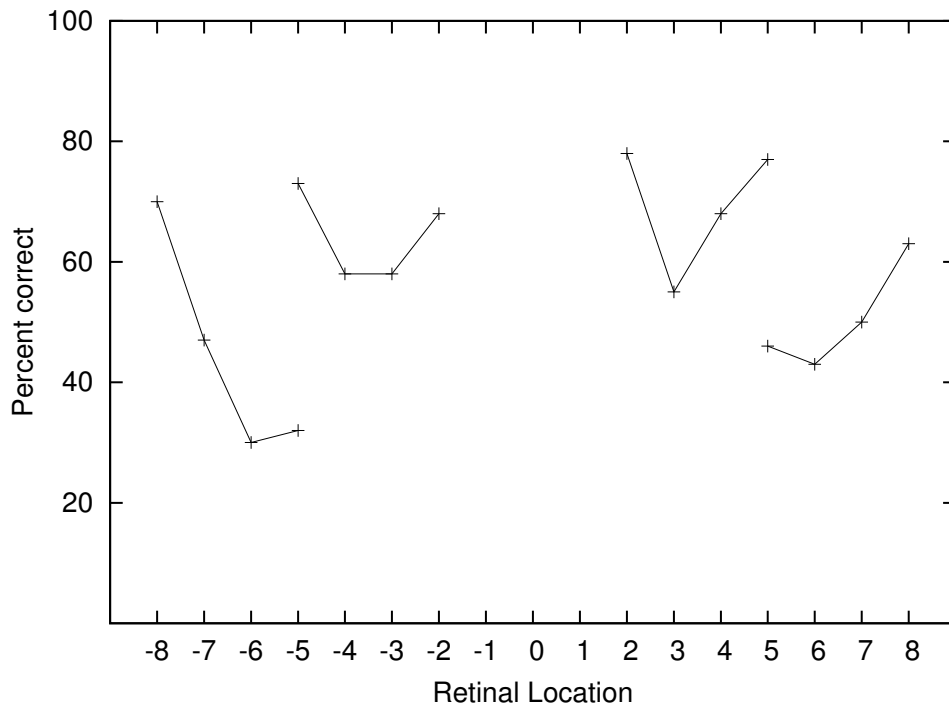


Figure 3:

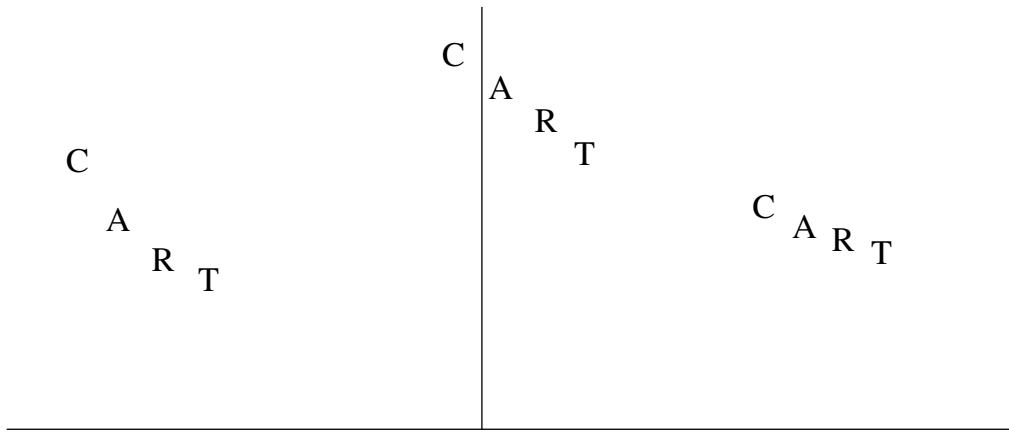


Figure 4:

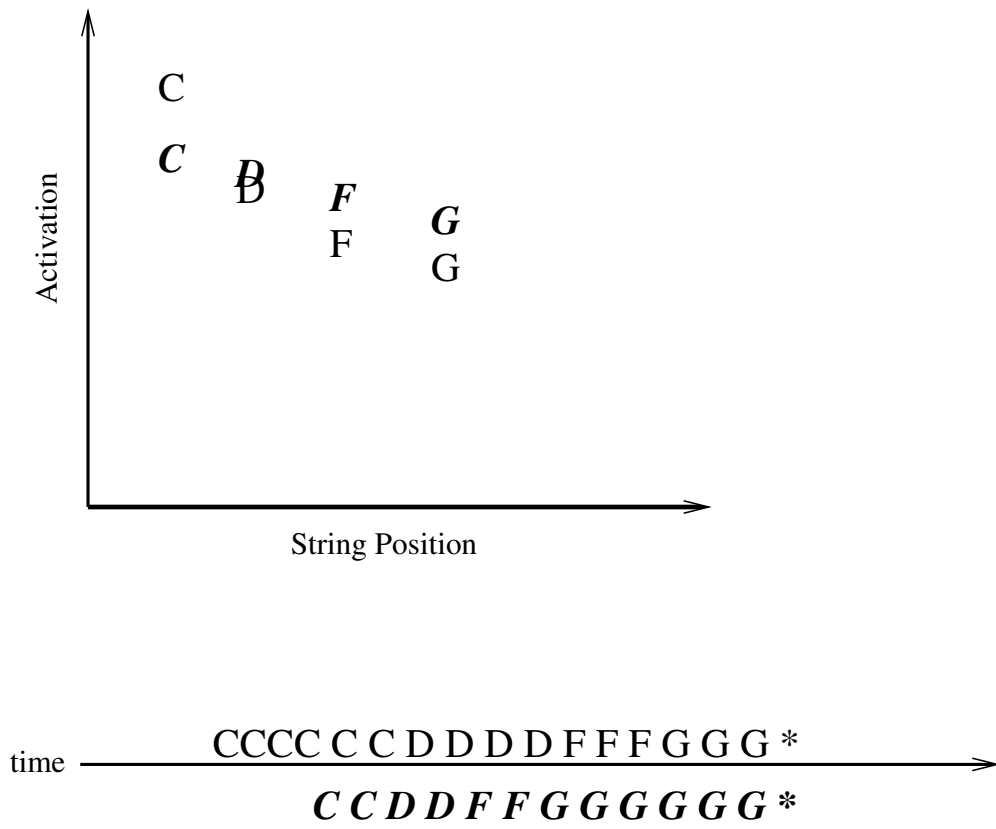


Figure 5:

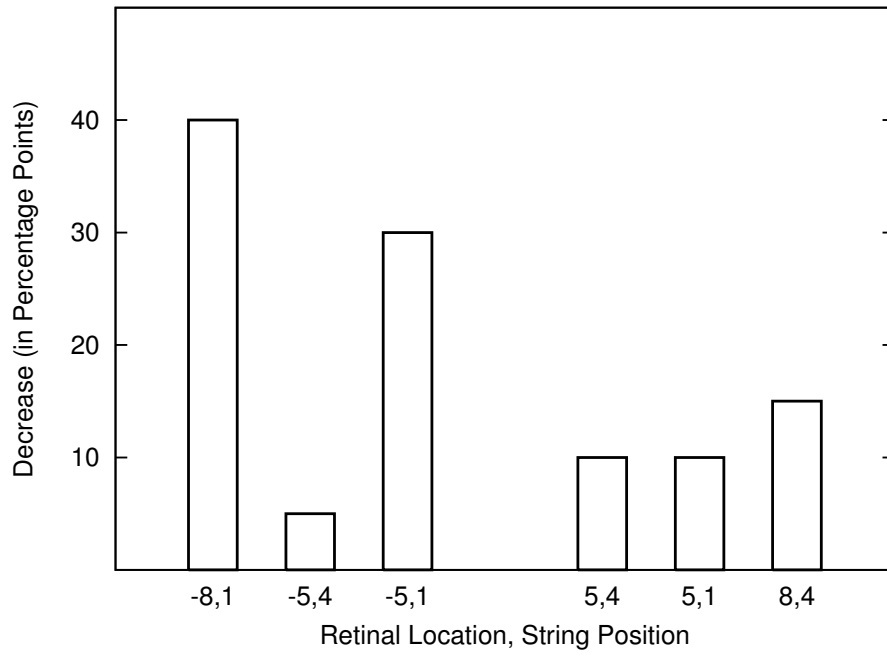


Figure 6:




Figure 7:




Figure 8:




Figure 9:




Figure 10: