

Comparison of the SERIOL and SOLAR Theories of Letter-Position Encoding

Carol Whitney

Department of Linguistics
University of Maryland
College Park, MD 20742
cwhitney@cs.umd.edu

Abstract

There has been increasing interest in the question of how the brain encodes the order of letters in a written word. This problem is of practical and theoretical interest, so it is important to distinguish between competing computational models. This article compares the SERIOL and SOLAR theories on their biological plausibility and ability to explain experimental results at the orthographic and lexical levels.

Key words: Visual Word Recognition, Letter-Position Encoding, Computational Modeling

1.0 Introduction

Following the publication of the SERIOL model of letter-position encoding (Whitney, 2001; Whitney & Berndt, 1999), there has been a growing interest in the issue of how the brain represents the order of letters in a written word, and a number of alternative models have been published (Davis & Bowers, 2006; Schoonbaert & Grainger, 2004; Dehaene, Cohen, Sigman, & Vinckier, 2005). The question of how letter position is encoded is a tractable problem of theoretical and practical importance. An understanding of orthographic processing may contribute to the treatment of developmental dyslexia, and to the identification of basic mechanisms in visual object recognition (Whitney & Cornelissen, 2005). Thus, it is crucial to differentiate between competing models.

The purpose of this article is to compare the SERIOL model to one particular alternative, the SOLAR model (Davis, 1999; 2007; Davis & Bowers, 2006; Davis, 2007). The structure of this article is as follows. In the next section, I discuss how to compare and evaluate neurocognitive models, both in general and with respect to letter-position encoding. I then review the SERIOL and SOLAR theories. Next, I evaluate the two theories according to the specified criteria, including a discussion of the problems with a comparison of the models presented in Davis and Bowers (2006).

2.0 On Models

In cognitive neuroscience, the term model can have two different meanings. It can refer to a theory of how information is represented and transformed by the brain during the performance of a task. Alternatively, it can denote a computer simulation of such processing. Clearly, a model in the former sense is more encompassing. An implementation requires the choice of specific values for all aspects of the computations being simulated, while a theory specifies the nature of the computations more abstractly. For clarity in the following, I will use theory to refer an abstract specification, and model to denote a simulation/implementation of a theory.

In understanding the brain, should we be more interested in theories or in simulations? Clearly, the goal is to specify the brain's computations in an abstract manner. That is, we are interested in computational theories. A simulation is merely an existence proof for a claim about a particular

aspect of a computational theory. That is, an implementation can demonstrate that some claim about a theory is true.

However, a simulation cannot demonstrate that a theory is unable to account for a certain pattern of results (unless a theory is very simple, making it tractable to simulate all possible combinations of parameters). Rather, to make such a claim, it is generally necessary to reason about the theory in a more abstract manner. This is trickier. A simulation gives concrete results, while reasoning about an abstract theory may seem nebulous. Yet, such an approach is necessary to establish that a theory in principle cannot explain certain data.

How then can we compare theories? First, a theory should explain relevant experimental data, and generate provable predictions. The strongest way to differentiate between two theories is to show that one accounts for the data, while the other cannot do so in principle. Note that evaluation of a specific implementation of a theory cannot invalidate that theory. There might be a different parameterization that would provide a much better fit to the data than the one being considered. After all, the parameters of a simulation are determined by fitting to existing data. If current parameters don't provide a good fit to new data, it may simply mean that they should be re-adjusted. Rather, to invalidate a theory it would be necessary to show that it cannot accommodate both the new and old data. Of course, a theory can be specified in varying levels of detail. It might be possible to show that some details are incorrect, but this would not falsify the overall theory. Thus, in order to experimentally differentiate between two theories, it is necessary to design experiments that address their core differences.

Another important dimension in comparing theories is biological plausibility. Does the theory explain how the computations could be performed at a neural level? Are the proposed mechanisms consistent with what is known about neural function? Or are the specified computations merely heuristics, with little connection to brain function?

How do these principles apply to the issue of letter-position encoding? First we need to define the problem. A string is initially represented in a retinotopic manner in the early visual areas. This must be transformed into an abstract, location-invariant encoding of letter order, which is matched

against stored orthographic word forms. Thus a complete theory of orthographic processing should specify the neural computations underlying the transformation of the retinotopic representation into a lexical encoding. Note that question of how letters are recognized is not directly addressed. Rather, given the ability to recognize letters, how is a spatial encoding transformed into an abstract representation of letter position, and how does this representation activate the lexical level?

Such a theory should explain behavioral data related to orthographic processing. Because a complete theory addresses the transformation from letters to words, this includes data both on lexical stimuli and on unpronounceable stimuli, where processing is presumably limited to the letter level. As discussed below, a variety of experiments have shown that lexical and orthographic phenomena interact with visual field of presentation and reading direction. Therefore, a theory of letter-position encoding should explain the effects of retinal location and processing direction.

Thus theories of letter-position encoding should be evaluated on their completeness (specification of the spatial-to-lexical transformation), their biological plausibility, and their inherent ability to explain experimental results at the letter and word levels, across presentation locations and languages. In the following, I review the SERIOL and SOLAR theories, and then compare them on these criteria.

3.0 Review

First, some terminology is specified. An italicized capitalization indicates a computational unit that represents the given item. For example, *CART* denotes a word unit that encodes the word CART. I use base word to denote the word encoded by the word unit under consideration, and comparison string to denote the word or string being compared to the base word. For example, we may be interested in how much the stimulus CAT activates CART. Here, CAT is the comparison string, and CART is the base word.

[Figure 1 about here]

In the SOLAR theory (right side of Figure 1), a sweep across a retinotopic encoding creates graded activations across abstract letter representations (Davis, 1999). That is, an abstract letter unit is

activated when the scan first falls on that letter, and its activation level increases until the end of the string is reached. As a result, the initial letter attains the highest activation, and activation levels decrease across the string. Learned weights on letter-to-word connections are proportional to the letter activation pattern.

Lexical activations are then computed as follows. For each word, the letter activations are compared to the weights to yield the difference between letter activations in the stimulus and the word. Then a symmetric function is centered around each difference. This function has a value of 1 at the center and decreases as distance from the center increases. The difference functions are summed across the letters, and the maximal value of this sum is divided by the number of letters in the word (Davis & Bowers, 2006).

Intuitively speaking, a word's activation level increases with the degree of overlap among the positional differences. For example, consider the comparison strings PAINTS and STAINS with respect to the base word WINTER, and assume that activation levels decrease across string positions from 6 to 1 in increments of 1. PAINTS yields positional differences of -1 for each of the shared letters I, N and T (e.g., the activation of N in WINTER is 3 and in PAINTS is 4; the difference is -1), while STAINS yields the following positional differences: -2 for I, -2 for N, and 2 for T. Because PAINTS yields more consistency among the positional differences than STAINS, the stimulus PAINTS would activate WINTER more than STAINS would.

In the SERIOL theory (left side of Figure 1) (Whitney & Berndt, 1999; Whitney, 2001; 2004a), learned left-to-right lateral inhibition (for a left-to-right language) at the feature level creates an activation gradient. This activation gradient interacts with letter units that undergo sub-threshold oscillations, creating sequential firing of these letter units. Open-bigram units respond to pairs of letters that fire in a particular order; such pairs are not necessarily contiguous within the stimulus. For example, the stimulus SPAT would activate *S*, then *P*, then *A*, then *T*. This firing pattern would activate open bigrams #*S*, *SP*, *PA*, *SA*, *AT*, *PT*, *ST*, and *T#*, where # represents a word boundary. Bigram activation levels decrease with increasing separation between the constituent letters (e.g., the activation of *SP* is higher than *SA* in this example). A maximum allowable separation of two letters is assumed (e.g., CRATE would not activate *CE*). Bigrams activate words via the standard

connectionist mechanism, where a word's activation level is given by the dot product of the bigram and weight vectors. That is, learned weights on bigram-to-word connections are proportional to the bigram activation pattern for each word; each bigram's activation is multiplied by the corresponding weight, and these products are summed to give a word's activation level.

Note the core differences between the two theories. In SOLAR, serial processing creates an activation gradient; in SERIOL, an activation gradient induces serial processing. The theories differ in the nature of the highest pre-lexical representation. In SOLAR, this level encodes the position of individual letters. In SERIOL, this level encodes relative position between letters, via ordered letter pairs. The lexical activation functions are also quite different. SERIOL uses the standard dot-product mechanism, while SOLAR specifies a complex function of activation differences.

Next I compare the theories. Both specify how a spatial encoding is transformed into a lexical encoding, meeting the completeness requirement. However, the theories differ in their ability to fulfill the other requirements.

4.0 Neural Plausibility

As discussed above, both theories assume serial processing. Such processing would have to be very rapid, on the order of milliseconds per letter, in order to create the appearance of parallel processing. Indeed, (English) experiments in which stimulus duration is varied have provided evidence for left-to-right processing of letters on the scale of tens of milliseconds per letter (Nice & Harcum, 1976).

SOLAR assumes a scan across the letters (Davis, 1999). How could such a scan occur? One possibility is serial allocation of attention across the string. However, top-down control of such sweep could not meet the stringent timing requirements; Wolfe, Alvarez & Horowitz (2000) showed that volitional shifts of attention proceed relatively slowly, on the time scale of 200 ms per shift. Furthermore, the ability to read single words is intact in patients with massive, bilateral damage to the parietal lobes, indicating that deployment of attention is not necessary for letter-position encoding (e.g., Vinckier et al., 2006). Rather, rapid serial processing would have to be performed in an automatic manner along the ventral visual stream, consistent with the fact that left

occipitotemporal damage results in slow, overtly serial processing of letter strings (i.e., letter-by-letter reading) (Cohen et al., 2004). That is, left occipitotemporal damage destroys the ability to process strings normally, while parietal damage does not. If normal visual word recognition depends on very rapid serial processing, this processing must be driven by left occipitotemporal cortex in a bottom-up manner, rather than by a top-down attentional scan from parietal cortex. However, SOLAR does not provide a mechanism for such an automatic serial encoding.

In contrast, SERIOL specifies the neural underpinnings of the serial processing. A top-down attention gradient is taken to drive the learning of left-to-right inhibition, to form a feature-level activation gradient (Whitney, 2004a; Whitney & Cornelissen, 2005). Based on neural mechanisms described by others (Hopfield, 1995; Lisman & Idiart, 1995), the gradient interacts with letter nodes that oscillate in excitability, resulting in a temporal encoding. The oscillatory mechanism is consistent with the increased occipital theta-band power observed in normal subjects during reading (Klimesch et al., 2001), and is taken to be an innate part of visual object recognition in general (Whitney, 2004a; Whitney & Cornelissen, 2005). Indeed, the particulars of the oscillatory mechanism predict a coupling between gamma power and theta phase, and such a coupling has recently been observed in an EEG experiment on visual perception (Demiralp et al., 2007). Thus the serial mechanism is well specified, biologically plausible, and operates rapidly in a bottom-up manner (after learning to read). As discussed below, these mechanisms also account for perceptibility patterns at the letter level, which SOLAR cannot explain.

The theories also differ in the plausibility of their lexical activation functions (arrow connecting highest two levels for each model in Figure 1). In neural modeling, a connection weight represents the efficacy of a synapse (or a set of synapses). The larger the connection weight, the better the transfer of the signal from the sending neuron(s) to the receiving neuron(s). This is modeled by multiplying the activation of the sending unit by the connection weight, to obtain the input to the receiving unit. Lexical activations in the SERIOL theory are based on these standard assumptions.

In contrast, the SOLAR function computes the difference between the activation of the sending unit and the connection weight. In this case, the connection weight does not model the efficacy of the transfer. As far as I know, no one has ever identified or proposed a neural mechanism that

directly computes the difference between a connection weight and an activity level. Rather, such a comparison of two quantities would entail instantiating them both as activity levels on different units, with the appropriate network connectivity. This is an inefficient mechanism for the parallel comparison of a stimulus. Essentially, each word unit would have to instantiate its connection weights as a spiking pattern. Such an approach would be metabolically expensive. In addition, the positional differences are then mapped onto functions, the functions are summated, and the maximal value is chosen. SOLAR does not specify how these computations would be carried out in a neural substrate.

In the standard dot-product model, activity simply flows from the sending neurons to the receiving neurons. Thus, in SERIOL, open bigrams directly activate lexical representations without any additional intermediate spiking activity. Moreover, the proposal that the highest pre-lexical representation encodes relationships between letter pairs is supported by a recent fMRI study, which showed that an area of left middle fusiform gyrus is uniquely sensitive to bigram probabilities (Binder et al., 2006). Indeed, the authors conclude (p. 742) “this region processes language-specific orthographic structure represented at the level of letter combinations, as has been proposed by several previous authors (Dehaene et al. 2005; Whitney, 2001).”

Thus, the SERIOL specification of neural activity is more complete and biologically plausible than the SOLAR specification. While SERIOL may seem more complex at first glance, this is only because SERIOL describes low-level mechanisms in detail, while SOLAR does not. At the lexical level, the SERIOL specification is actually much simpler than SOLAR. Next I consider the ability of the theories to explain experimental results at the letter and word levels.

5.0 Letter Level

Most theories of letter-position encoding have focused on phenomena at the lexical level, such as the results of form-priming experiments. However, non-word strings produce distinctive perceptual patterns. Arrays of alphanumeric characters produce an external-character advantage (i.e. the first and final characters are perceived better than the internal characters). However, arrays of non-alphanumeric symbols do not; the first and last symbols of a centrally fixated array are the least well perceived, as shown by two independent studies (Hammond & Green, 1982; Mason,

1982). Thus the external-letter advantage cannot be due to a general lack of low-level lateral inhibition (because it is not present for non-alphanumeric symbols), but rather must arise from string-processing mechanisms.

Furthermore, there are asymmetries in the perceptibility patterns for right versus left visual-field presentation (LVF vs. RVF), which are modulated by reading direction. For example, consider trigram identification in English. In the LVF/RH, the first letter is perceived much better than the third letter; in the RVF/LH, the first and third letters are perceived equally well (Hellige, Cowin & Eng, 1995). However, for languages read from right-to-left, this pattern reverses. In the RVF/LH the first letter is perceived better than the third; in the LVF/RH, the first and third letters are perceived more equally well (Adamson & Hellige, 2006; Eviatar, 1999). Nazir et al. (2004) have also shown that hemifield perceptibility patterns vary with reading direction. Such results could not stem from purely perceptual processes, nor from innate hemispheric specializations.

These experimental data clearly show that perceptibility patterns in non-word strings are influenced by the demands of representing strings for reading. Therefore, any theory of letter-position encoding should explain how and why these patterns emerge. SOLAR allows stronger weights on exterior than interior letters, but doesn't explain how an external-letter superiority comes about. (As discussed above, it couldn't arise from low-level perceptual processes.) In contrast, SERIOL explains why the external-letter superiority arises, and how the patterns for external and internal letters vary with visual field, reading direction, and exposure duration. These accounts are based on SERIOL's specification of hemispheric specialization in direction-specific inhibition at the feature level, coupled with the oscillatory mechanism that produces serial firing at the letter level. See Whitney (2001; 2004a) and Whitney and Cornelissen (2005) for details.

Moreover, this analysis has led to novel experimental results. It suggested that VF asymmetries in lexical decision also stem from orthographic activation patterns (Whitney, 2004b), predicting that specific positional adjustments of contrast level should reverse the normal patterns of asymmetry (i.e., make LVF presentation give the usual RVF results, and vice versa). These predictions were experimentally confirmed (Whitney & Lavidor, 2004; Whitney & Lavidor, 2005), for asymmetries related to string length (Young & Ellis, 1985) and neighborhood size (Lavidor & Ellis, 2002). No

other existing model of letter-position encoding could have generated such predictions, nor can explain these results. Thus these results are highly specific to the SERIOL theory, and provide strong evidence for proposed processing at the feature and letter levels.

It is also of interest to note that the source of the asymmetry of the length effect has been a topic of debate for decades. Whitney and Lavidor (2004) showed, for the first time, how to abolish the normal LVF length effect, establishing that it does not arise from a hemisphere-specific mode of processing (Young & Ellis, 1985).

6.0 Word level

Davis and colleagues have performed several experiments giving results that are consistent with the SOLAR theory, but not specific to it (Bowers, Davis & Hanley, 2005; Davis & Bowers, 2004; Davis & Taft, 2005). These experiments demonstrated general evidence against position-specific letter codes, and these data can be explained equally well by the SERIOL and SOLAR theories. In contrast, Davis and Bowers (2006) presented experimental data that they claimed provided evidence against the SERIOL theory. I first consider the problems with this article, and then discuss how to differentiate between the theories.

6.1 Davis and Bowers (2006)

Davis and Bowers (2006) sought to differentiate the SOLAR and SERIOL theories based on match scores from implemented models. (A match score is the bottom-up input to the base word's word-unit for a given comparison string.) However, such an approach should not be used to make general statements about the comparative veracity of two theories, as discussed above.

Davis and Bowers focused on two types of comparison strings, which both contained one mismatching letter with respect to a base word. In one type, the matching letters all occurred in the same position as the base word (orthographic neighbor). In the other type, one of the internal matching letters was shifted by one position (Neighbor Once Removed - NIR). For example, consider the base word SHOP. STOP is an orthographic neighbor; the matching letters S, O, and P are in the same string position in both words. SOAP is an NIR of SHOP; the S and P are in the same position, while the O is shifted by one position.

[Figure 2 about here.]

For clarity, comparison strings will be specified in terms of the base word's letters. In our example, SOAP is represented 13x4, indicating that its first and fourth letters match the base word SHOP in identity and position, its second letter is the base word's third letter, and its third letter is not present in the base word. (See Figure 2 for an illustration of how a comparison string activates the base word.) The critical issue is whether neighbors (e.g., 1x34) are more similar than N1Rs (e.g., 13x4) to the base word.

Under the SOLAR theory, a neighbor will generate a higher match score than an N1R under any parameter setting, because there is more overlap among positional differences (i.e., all positional differences are 0 for neighbors, while this is not the case for N1Rs). Under the SERIOL theory, the ordering of match scores depends on the particular form of the bigram activation function. In the original specification (Whitney & Berndt, 1999), bigram activations varied with string position. Under this parameterization, the ordering depends on the original position of the mismatching letter: the neighbor 12x4 scores higher than the N1R 1x24, but the neighbor 1x34 scores lower than the N1R 13x4. That is, when the third letter of the base word is not included in the comparison string, neighbors score better than N1Rs; when the second letter is not included, N1Rs score better than neighbors. However, this positional assumption was dropped in Whitney (2004a). (The phenomena that were formerly explained by positional variations are now explained directly by seriality, so this change did not reduce explanatory capacity.) Under the current assumption that bigram activations depend only on the separation between the constituent letters, neighbors always score higher than N1Rs.

In a series of experiments using four- and five-letter words, Davis and Bowers (2006) found that neighbors are more similar than N1Rs to the base word. In comparing SOLAR and SERIOL, the authors based their analysis of SERIOL on the parameters specified in Whitney & Berndt (1999), and only considered match scores for conditions like 1x34 vs. 13x4, where N1Rs score better than neighbors. They mistakenly concluded that the SERIOL theory predicts an advantage for N1Rs over neighbors, and interpreted their results as evidence against SERIOL.

In their Discussion, Davis and Bowers did consider whether dropping the positional assumption would lead SERIOL to predict a different pattern of results. They conclude that the SERIOL theory is, in principle, not capable of generating the observed ordering. This is clearly incorrect. Consider, for example, neighbor 1x345 versus N1R 13x45 under the current parameters (Whitney, submitted), where an edge or contiguous bigram gives an activation of 1.0, a one-letter separation gives 0.8, a two-letter separation gives 0.4. Both comparison strings match on #1, 1~4, 45, and 5#, giving a total of 3.16 for the common matches. For 1x345, the remaining matches (1~3, 34, and 3~5) give:

$$(0.8 * 0.8) + (1.0 * 1.0) + (0.8 * 0.8) = 2.28.$$

For 13x45, the remaining matches (13, 3~4, and 3~5) give:

$$(1.0 * 0.8) + (0.8 * 1.0) + (0.4 * 0.8) = 1.92.$$

Thus, the raw scores are 5.44 for 1x345 versus 5.08 for 13x45; neighbor 1x345 gives a score that is 7.1% larger than N1R 13x45. How does the relative size of this effect compare with that given by the current SOLAR model (Davis, 2007)? The raw scores are 4.00 versus 3.72, giving an effect size of 7.5%. Thus, for five-letter base words, the current SERIOL model gives an effect in the observed direction that is of the same magnitude as the current SOLAR model.

For four-letter stimuli, the current SERIOL parameters also give a higher match score for neighbors, but the effect size is considerably smaller. This may indicate that the parameters need to be adjusted. Decreasing activation levels for edges and one-letter separations would increase the effect size. For example, if the activations for edge bigrams, and one- and two-letter separations are all set to 0.5, the effect size becomes 5.8%. However, the experimental results on four-letter stimuli are questionable. That study used the illusory-word paradigm, where two strings are presented and one is identified. Therefore, the four-letter results could reflect conscious, strategic effects that are outside the realm of normal orthographic processing. In contrast, the five-letter experiments used the form-priming method, which is known to tap exclusively into automatic orthographic processing (Grainger et al., 2006). It remains to be seen whether an advantage for neighbors would arise for four-letter stimuli under the automatic conditions of the priming paradigm.

There is also a potential problem with the five-letter results. In these experiments, both the prime

and the target were centered, and there was an effort to minimize effects of retinotopic overlap by presenting primes in 12-pt lower-case, and targets in 16-pt upper-case. However, many lower- and upper-case letters are quite similar, and a discrepancy in font-size has a minimal effect at fixation, because the prime and target letter are both centered at the same retinal location. The N1R conditions (13x45 and 12x35) mismatched the base word on the central letter, while the neighbor conditions (1x345 and 123x5) matched on the central letter. Thus the advantage for neighbors may simply reflect retinotopic overlap in the central position. While this overlap was not exact, an approximate match would still be sufficient to activate the same cortical feature detectors when the upper- and lower- case letters shared similar features. To clearly demonstrate an advantage for neighbors, the prime should be shifted relative to the target so that neighbors do not have a possible retinotopic advantage. So although both models predict a small advantage for neighbors over N1Rs for five-letter stimuli, it is unclear whether a difference of this size should necessarily be experimentally detectable; it may be the case that a better controlled experiment would show no difference between the two conditions, which would also be consistent with both models, under the assumption that the predicted difference is too small to detect.

Thus, it is unclear whether there truly is a strong advantage for neighbors over N1Rs. If future priming experiments were to convincingly demonstrate such an advantage for five-letter words but not four-letter words, the current SERIOL parameters would explain these results as well as the current SOLAR model does. If a neighbor advantage were also demonstrated for four-letter words, a readjustment of the SERIOL parameters may then be in order.

In sum, both the SOLAR and SERIOL theories can account for a greater similarity of neighbors over N1Rs. This advantage comes about for different reasons in the two models, which are related to how the models use positional information. In SERIOL, absolute string position is abstracted away; relative position is represented by open bigrams. An advantage for neighbors arises because the distance between letters is encoded in bigram activation levels. In SOLAR, an advantage for neighbors arises because there is more overlap in activation differences at the level of individual letters. Experiments comparing N1Rs and neighbors cannot differentiate between the underlying theories. In the following section, a comparison that can differentiate between the theories is discussed. First, however, we consider the implications of a recent modification to SOLAR .

The SOLAR theory has now been generalized to include parameterization of the function specifying the activation gradient (Davis, 2007). Originally, the activation drop-off across adjacent positions did not vary with string length (Davis & Bowers, 2006). Under this former assumption, the difference in the activations of letters 1 and 9 in 123456789 is much larger than the difference between letters 1 and 4 in 1234, for example. Under the current parameters, the drop-off between adjacent positions decreases with increasing string length such that the total drop-off across the first and last letters is approximately normalized. For example, the activation difference between 1 and 4 in 1234 is now about the same as the difference between 1 and 9 in 123456789. This parameterization provides better explanatory capacity in some cases where the comparison string is shorter than the base word (discussed in more detail below).

This new assumption essentially has the effect of “stretching out” a short comparison string with respect to a longer base word. As a result, match scores for short comparison strings are higher when the retained letters are sampled more evenly across the base word. For example, 1368 scores better than 1234 for a nine-letter base word (0.42 versus 0.31). However, one of the main points of Davis and Bowers (2006) was to claim that there is a strong advantage for maintaining the contiguity of the letters in the base word, and to argue that the SERIOL model inherently does not encode contiguity to a sufficient degree. They claim “the mechanism that matches input codes against previously learned codes must be sensitive to incongruities in letter contiguity” (Davis & Bowers, p. 550). For example, neighbor 12x4 was taken to have an advantage over N1R 1x24 because the contiguity of 1 and 2 is maintained in the neighbor, but not in the N1R. By this logic, 1234 should certainly score better than 1368, because 1234 replicates the contiguity of the base word, while 1368 does not. Indeed, 1234 did indeed attain a much higher score than 1368 under the SOLAR parameterization in Davis & Bowers (2006). However, the current SOLAR model actually makes the opposite prediction; it predicts a *cost* of contiguity when the comparison string is shorter than the base word.

Thus normalization of the activation gradient leads to a direct contradiction of the claims of Davis and Bowers (2006). However, the original absence of normalization caused other problems. For example, consider a six-letter base word, and the primes 1346 and 1-34-6, where the dashes are

literals. Note that the letters in 1-34-6 maintain the string positions of the base word, while the letters in 1346 do not. Without gradient normalization, SOLAR predicts that 1346 should provide significantly less facilitation than 1-34-6, because the letters in 1346 are shifted by varying amounts (with respect to the base word) while positional shifts in 1-34-6 are all equal (to 0). That is, there should be an advantage for maintaining absolute string position because it increases consistency in positional differences. However, numerous experiments have shown no advantage whatsoever for primes like 1-34-6 versus 1346 (Peresotti & Grainger, 1999; Grainger et al., 2006). If instead the SOLAR activation gradient is normalized, 1346 is essentially stretched out to become equivalent to 1-34-6, correctly giving no advantage. In contrast, SERIOL inherently predicts no advantage for maintaining string position, while at the same time predicting an advantage for 1234 over 1368.

6.2 Differentiating between the Theories

For many stimuli, the SERIOL and SOLAR theories give similar predictions because the open-bigram computation and the activation-difference function yield comparable results. When letter order is violated, fewer open-bigrams are activated and variance in positional shifts increases, yielding decreased lexical activation under both theories. However, it is possible to differentiate between the two proposals using stimuli carefully chosen to exploit their differences. Consider a seven-letter base word, and the primes 1237 and 1654327. Both match on the external letters, so edge bigrams and edge effects are equivalent for each and will be ignored.

Under the SERIOL theory, 1237 must score better than 1654327. 1237 activates open-bigrams 12, 23, and 1~3, while 1654327 activates open-bigrams 1~4 and 4~7. Thus 1237 activates more bigrams with larger weights to higher levels, giving it a strong advantage over 1654327. For example, current parameters give a normalized match score of 0.39 for 1237 versus 0.20 for 1654327.

In contrast, 1654327 scores better than 1237 under the SOLAR theory. For 1654327, letters 1, 4, and 7 match the base word's position, each contributing 1.0 to the raw score. Letters 3 and 5 are shifted from their original positions, and each contributes significantly less than 1.0. (2 and 6 are shifted too far to contribute anything.) For 1237, 1, 2, and 3 each contribute around 1.0, while 7 is shifted and contributes significantly less than 1.0. (Depending on the parameters, 1, 2, and 3 may

each contribute somewhat less than 1.0, but this shortfall can be thought of as reducing the contribution from 7.) Thus 1654327 and 1237 both have three maximal contributors, while the former has two partial contributors and the latter only has one.

The exact values of the partial contributions in each condition will depend on the SOLAR parameters, but for any given set of parameters, the contribution from 3 (or 5) in 1654327 is at least as large as the contribution from 7 in 1237. Intuitively, this can be seen by noting that the 7 in 1237 is shifted a distance of three positions with respect to the base word, while the 5 or 3 in 1654327 is shifted a distance of only two positions. A smaller shift will tend to increase overlap and yield a larger contribution. (I acknowledge that this analysis is a simplification of the SOLAR theory, and the details are not precisely correct. However, it gives a reasonable approximation of the outcome of the complex computations, for any parameterization of the activation gradient.) Therefore 1654327 will score higher than 1237 because the sum of the partial contributions is larger. For example under the current SOLAR parameters (Davis, 2007), 1654327 generates a normalized match score of 0.53, and 1237 gives 0.48. The key point is that the prime 1654327 must provide *at least as much* facilitation as 1237 under SOLAR. In contrast, SERIOL predicts that 1654327 should provide significantly less facilitation than 1237.

Thus the two theories give different predictions for these conditions. Studies have previously shown that facilitation occurs when the target's letter order is preserved, but not when it is violated. For example for seven-letter targets, primes 1357 and 13457 speed reaction times (relative to a same-length control prime of mismatching letters), while 1537 and 15437 do not provide any facilitation (Grainger et al., 2006). Generalizing from these results, it is likely that 1237 would provide facilitation and 1654327 would not, contrary to the prediction of the SOLAR theory. Of course, data from these specific primes are necessary to definitively decide the matter; the primary purpose of this discussion is to demonstrate that it is indeed possible to design priming conditions that capture *inherent* differences between the models.

In sum, SOLAR seeks to achieve the relative-order constraint without using multi-letter units. However, the resulting activation function is quite elaborate and does not achieve this goal in some cases. It is difficult to see any possible advantage for eschewing multi-letter units. Higher-order

units yield a simpler, more robust, and more biologically-plausible theory. As discussed above, an fMRI study led the authors to conclude that orthographic encoding involves a level of multi-letter units, as in SERIOL (Binder et al., 2006). Furthermore, a recent single-cell study of macaque inferotemporal cortex revealed neurons that are selective for a particular spatial configuration of two features (Brincat & Connor, 2006). Similarly, I assume that cells in human inferotemporal cortex become tuned to specific arrangements of two letters (i.e., become open-bigram detectors).

7.0 Conclusion

The computations in the SERIOL theory are biologically plausible, can account for the experimental results of Bowers and Davis (2006), explain letter perceptibility patterns, and have led to novel, confirmed predictions that are highly specific to the theory. The computations in the SOLAR theory are implausible, provide no explanatory advantage, do not explain letter perceptibility patterns, and have not generated any experimental results that are specific to the theory. Thus, an up-to-date and thorough comparison of the SOLAR and SERIOL theories shows the advantages of SERIOL and the disadvantages of SOLAR.

References

- Adamson, M.M. & Hellige, J.B (2006) Hemispheric differences for identification of words and nonwords in Urdu-English bilinguals. *Neuropsychology*, 20, 232-248.
- Binder, J.R., Medler, D.A., Westbury, C.F., Liebenthal, E. & Buchanan L. (2006) Tuning of the human left fusiform gyrus to sublexical orthographic structure. *Neuroimage*. 33, 739-748.
- Brincat, S.L. & Connor, C.E. (2006) Dynamic shape synthesis in posterior inferotemporal cortex. *Neuron*, 49, 17-24.
- Bowers, J. S., Davis, C. J., & Hanley, D. A. (2005a). Automatic semantic activation of embedded words: Is there a 'hat' in 'that'? *Journal of Memory and Language*, 52, 131-143.
- Cohen, L., Henry, C., Dehaene, S., Martinaud, O., Lehericy, S., Lemer, C., & Ferrieux S. (2004) The pathophysiology of letter-by-letter reading. *Neuropsychologia*. 42, 1768-1780.
- Davis, C. (1999) The Self-Organizing Lexical Acquisition and Recognition (SOLAR) Model of Visual Word Recognition. Doctoral Dissertation. University of New South Wales.
- Davis, C. J. (2007) MatchCalculator software, available at www.pc.rhul.ac.uk/staff/c.davis/Utilities/
- Davis, C.J. & Bowers, S.J. (2004) What do letter migration errors reveal about letter position coding in visual word recognition? *Journal of Experimental Psychology: Human Perception and Performance*, 30, 923-941.
- Davis, C.J. & Bowers, S.J. (2006) Contrasting five different theories of letter-position coding: Evidence from orthographic similarity effects. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 535-557.

Davis, C. J. & Taft, M. (2005). More words in the neighborhood: Interference in lexical decision due to deletion neighbors. *Psychonomic Bulletin & Review*, 12, 904-910.

Dehaene, S., Cohen, L., Sigman, M. & Vinckier, F. (2005) The neural code for written words: a proposal. *Trends in Cognitive Science*. 9, 335-341.

Demiralp, T., Bayraktaroglu, Z., Lenz, D., Junge, S., Busch, N.A., Maess, B., Ergen, M. & Herrmann, C.S. (2007) Gamma amplitudes are coupled to theta phase in human EEG during visual perception. *International Journal of Psychophysiology*, in press.

Eviatar, Z. (1999) Cross-language tests of hemispheric strategies in reading nonwords. *Neuropsychology*, 13, 498-515.

Hammond, E.J. & Green, D. (1982) Detecting targets in letter and non-letter arrays. *Canadian Journal of Psychology*, 36, 67-82.

Hellige, J.B., Cowen, E.L., & Eng, T.L. (1995) Recognition of CVC syllables from LVF, RVF, and central locations: Hemispheric differences and interhemispheric interactions. *Journal of Cognitive Neuroscience*, 7, 258-266.

Hopfield, J.J. (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376, 33-36.

Lavidor, M & Ellis, A. (2002) Word length and orthographic neighborhood size effects in the left and right cerebral hemispheres. *Brain and Language*, 80, 45-62.

Lisman, J.E. & Idiart, M.A.P. (1995) Storage of 7 (2 short-term memories in oscillatory subcycles. *Science*, 267, 1512-1515.

Mason, M. (1982) Recognition time for letters and nonletters: Effects of serial position, array size,

and processing order. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 724-738.

Nice, D.S. & Harcum, E.R. (1976) Evidence from mutual masking for serial processing of tachistoscopic letter patterns. *Perceptual and Motor Skills*, 42, 991-1003.

Nazir, T.A., Ben-Boutayab, N., Decoppet, N., Deutsch, A. & Frost, R. (2004) Reading habits, perceptual learning, and recognition of printed words. *Brain and Language*, 88, 294-311.

Peressotti, F. & Grainger, J. (1999) The role of letter identity and letter position orthographic priming. *Perception & Psychophysics*, 61, 691-706.

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.

Vinckier, F., Naccache, L., Papeix, C., Forget, J., Hahn-Barma V., Dehaene, S. & Cohen, L. (2006) "What" and "where" in word reading: Ventral coding of written words revealed by parietal atrophy. *Journal of Cognitive Neuroscience*, 18, 1998-2012.

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. (2004a) Investigations into the Neural Basis of Structured Representations. Doctoral Dissertation. University of Maryland.

Whitney, C. (2004b) Hemisphere-specific effects in word recognition do not require hemisphere-specific modes of access. *Brain and Language*, 88, 279-293.

Whitney, C. (submitted) Supporting the serial in the SERIOL model.

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. & 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.

Wolfe, J.M., Alvarez, G.A. & Horowitz, T.S. (2000) Attention is fast but volition is slow. *Nature*, 406, 691 - 691.

Young, A.W. & Ellis, A.W., (1985) Different methods of lexical access for words presented in the left and right visual hemifields. *Brain and Language*, 24, 326-358.

Figure Captions

Figure 1: Comparison of the SERIOL and SOLAR models for a language read from left to right. Each box specifies a level of representation, where *spatial* denotes a retinotopic encoding and *abstract* denotes a non-retinotopic encoding. Additional information about the nature of the representation is given in italics. The text near the arrows specifies the proposed mechanisms underlying the transformations between levels.

Figure 2: Illustration of the activation of BLANKET by a comparison string of the form 1347 (i.e., the string BANT). Letter units fire serially, activating open-bigram units. The activation level of an open-bigram unit depends on the amount of time between the firing of the constituent letter units. Each connection weight records the corresponding activation for the base word itself. (Actually, weights also depend on the length of the base word, but this is ignored for simplicity.) Open-bigrams AT and BT have weights of 0 (not shown) because the constituent letters are too far apart in the base word to activate these open-bigrams. Each open-bigram unit generates output equaling the product of its activation and weight. The input to the word unit is the sum of these products, giving the raw score. The normalized score is the raw score divided by the raw score for an exact match (i.e., for the base word itself). The normalized score here is $4.52 / 11.84 = 0.38$.




Figure 1




Figure 2