



Published in final edited form as:

Psychol Rev. 2012 April ; 119(2): 417–430. doi:10.1037/a0027175.

Competition and cooperation among similar representations: Toward a unified account of facilitative and inhibitory effects of lexical neighbors

Qi Chen and Daniel Mirman

Moss Rehabilitation Research Institute

Abstract

One of the core principles of how the mind works is the graded, parallel activation of multiple related or similar representations. Parallel activation of multiple representations has been particularly important in the development of theories and models of language processing, where co-activated representations (“neighbors”) have been shown to exhibit both facilitative and inhibitory effects on word recognition and production. Researchers generally ascribe these effects to interactive activation and competition, but there is no unified explanation for why the effects are facilitative in some cases and inhibitory in others. We present a series of simulations of a simple domain-general interactive activation and competition model that is broadly consistent with more specialized domain-specific models of lexical processing. The results showed that interactive activation and competition can indeed account for the complex pattern of reversals. Critically, the simulations revealed a core computational principle that determines whether neighbor effects will be facilitative or inhibitory: strongly active neighbors exert a net inhibitory effect and weakly active neighbors exert a net facilitative effect.

Keywords

Interactive activation; competition; neighbor effects; neighborhood effects; lexical processing

Introduction

Across theoretical frameworks and domains of cognition, one of the core principles of how the mind works is the graded, parallel activation of multiple related or similar representations. This idea is a central tenet of parallel distributed processing (PDP) models of cognition (e.g., McClelland, 1993; Rumelhart, McClelland, and the PDP Research Group, 1986), as well as exemplar models of memory (e.g., Kalish, Lewandowski, & Kruschke, 2004; Medin & Schaffer, 1978), and Bayesian approaches to cognition (e.g., Griffiths, Kemp, & Tenenbaum, 2008). Parallel activation of multiple representations has been particularly important in the development of theories and models of language processing, where these related representations have been called “neighbors” and studied extensively in a wide range of tasks and contexts. Although there is broad agreement on the principle of

Correspondence: Qi Chen, Moss Rehabilitation Research Institute, 50 Township Line Rd., Elkins Park, PA 19027, Phone: 1 (215) 456 9585, chenqi.research@gmail.com.

Publisher's Disclaimer: The following manuscript is the final accepted manuscript. It has not been subjected to the final copyediting, fact-checking, and proofreading required for formal publication. It is not the definitive, publisher-authenticated version. The American Psychological Association and its Council of Editors disclaim any responsibility or liabilities for errors or omissions of this manuscript version, any version derived from this manuscript by NIH, or other third parties. The published version is available at www.apa.org/pubs/journals/rev

parallel activation, the consequences are quite varied and, to date, there is no formal unified account of why lexical neighbors exert facilitative effects in some contexts and inhibitory effects in others. This is because research on neighbor effects has been almost entirely isolated by domain: there are sophisticated models of, for example, reading aloud (e.g., DRC: Coltheart et al., 2001; CDP+: Perry, Ziegler, & Zorzi, 2007; SOLAR: Davis, 2010), which capture neighbor effects in reading aloud in great detail, but do not address neighbor effects in picture naming or spoken word recognition or other tasks. Here, we take an orthogonal approach: instead of building a detailed model of one task, we build a simple and general model designed to be applicable across domains and use it to uncover the computational principle that underlies the contrasting results in the literature. A corollary of this approach is that our model is not meant to compete with existing models of lexical processing or neighborhood effects; rather, it is meant to be a step toward bridging across existing models by identifying underlying computational principles that would need to hold in a full cross-domain model of lexical processing.

We first briefly review the literature on lexical neighbor effects, documenting which contexts elicit facilitative effects and which contexts elicit inhibitory effects. We then describe a simple interactive activation and competition (IAC; e.g., McClelland & Rumelhart, 1981) framework for exploring these effects. Many researchers have used IAC frameworks to explain lexical neighborhood effects, either by intuition or through implemented computational models, but there has been little effort to explain why the same framework would predict inhibitory neighbor effects in some cases and facilitative neighbor effects in other cases. Using simulations, we show that these contrasts can arise in a single computational framework and provide a unified account of why some neighbors facilitate and others inhibit word processing.

Visual word recognition

Facilitative effects of orthographic and phonological neighbors

One of the first demonstrations of neighborhood effects in language processing was the finding that printed words with many orthographic neighbors were recognized more quickly than words with few neighbors (Andrews, 1989, 1992; Forster & Shen, 1996; Johnson & Pugh, 1994; Sears, Hino, & Lupker, 1995; for a review see Andrews, 1997). Most of these studies used a lexical decision or naming (reading aloud) task, though the latter task also involves word production processes that, as we will review, are also sensitive to neighborhood effects. The typical definition of an orthographic neighborhood is “Coltheart’s N” (Coltheart, Davelaar, Jonasson, & Besner, 1977), which is the number of words that can be created by changing a single letter of a target word. For example, the target word *mint* has neighbors such as *mitt*, *tint*, *lint*, and *pint*. Other researchers have proposed alternative measures of orthographic neighborhood (e.g., Yarkoni, Balota, & Yap, 2008) but all of the measures have retained the core notion of letter-based similarity and have demonstrated facilitative effects of orthographic neighbors.¹

Phonological neighbors – words that can be created by changing a single phoneme – can be somewhat different from orthographic neighbors. For example, the phonological neighborhood of *mint* includes orthographic neighbors such as *tint* and *lint*, as well as additional words such as *mince*, *meant*, and *minnow*, and excludes some orthographic

¹The facilitative effects of neighbors on visual word recognition appear to be stronger for low frequency words than high frequency words (e.g., Coltheart et al., 2001; Yarkoni et al., 2008; Davis, 2010, makes the even stronger claim that there are no effects of neighbors on high frequency words). Like the well-known frequency-by-regularity interaction, this could arise simply because recognition of high frequency words is fast, making it difficult to detect a neighbor facilitation effect (for related discussion see, e.g., Plaut, McClelland, Seidenberg, & Patterson, 1996; and for a similar account in a very different domain see discussion of the asymmetry of word and ink color interference effects in the Stroop task in Cohen, Dunbar, & McClelland, 1990).

neighbors, such as *pint*. Nevertheless, phonological neighbors also exert facilitative effects on visual word recognition (Yates, 2005; Yates, Locker, & Simpson, 2004).

Inhibitory effects of higher-frequency neighbors

The studies reviewed above suggest cooperation among lexical neighbors during visual word recognition. Since word frequency facilitates word recognition, one might imagine that neighbors that are more frequent would be more facilitative. The data reveal the opposite pattern: when neighbors are more frequent than the target word, they exert an inhibitory effect on target word recognition (Davis, Perea, & Acha, 2009; Ferraro & Hansen, 2002; Grainger, 1990; Grainger & Jacobs, 1996; Grainger, O'Regan, Jacobs, & Segui, 1989, 1992; Grainger & Segui, 1990).

These results indicate that there must be some balance between facilitative and inhibitory effects of orthographic neighbors. The need for this kind of balance is further demonstrated by inhibitory effects of transposed letter neighbors: words that have a transposed letter neighbor, such as *salt* – *slat*, are recognized more slowly than matched words that do not have a transposed letter neighbor, such as *halt* (Acha & Perea, 2008; Andrews, 1996; Johnson, 2009). Andrews (1996) showed that both the facilitative effect of letter substitution neighbors and the inhibitory effect of transposed letter neighbors can be demonstrated within a single data set and concluded that “It remains to be seen whether it is possible to find a single set of parameters that allow successful simulation of both phenomena.” The search for models and parameters that can account for these contrasting patterns has been a major theme in the development of theories and models of visual word processing (for a review, see Grainger, 2008).

Spoken word recognition: Inhibitory effects of neighbors

When words are presented auditorily, rather than visually, the effect of neighbors reverses to inhibit word recognition (e.g., Luce, 1986; Luce & Pisoni, 1998). Indeed, neighborhood probability – a measure that combines relative word frequency and number of neighbors – accounts for about 15% of the variance in tasks like lexical decision and word repetition (Luce, 1986; Luce & Pisoni, 1998). The next best predictor is frequency alone, which only accounts for about 5% of the variance. The inhibitory effect of phonological neighbors on spoken word recognition has also been shown in word-to-picture matching (e.g., Magnuson, Dixon, Tanenhaus, & Aslin, 2007), gating (e.g., Garlock, Walley, & Metsala, 2001), priming (Dufour & Peereman, 2003a, 2003b; Goldinger, Luce, & Pisoni, 1989) and other paradigms and using different definitions of phonological neighbors (e.g., Luce & Pisoni, 1998; Magnuson et al., 2007; see also Benki, 2003).

Although most evidence shows inhibitory effects of lexical neighbors on spoken word recognition, partial activation of many lexical representations may have facilitative effects at sub-lexical levels of speech processing. The best evidence of such facilitative top-down effects is a bias to identify ambiguous phonemes towards a denser lexical neighborhood (Newman, Sawusch, & Luce, 1997; see also Boyczuk & Baum, 1999). For example, an ambiguous sound between /g/ and /k/ was more likely to be identified as /g/ when followed by *-ice*, presumably because *gice* has higher neighborhood density than *kice*. Sub-lexical facilitation has also been implicated in facilitative effects of phonotactic probability – the relative likelihood of phoneme pairs (e.g., Vitevitch & Luce, 1998, 1999; Vitevitch, Luce, Pisoni, & Auer, 1999; see also Luce & Large, 2001). However, other researchers have questioned this result based on failures to replicate (Lipinski & Gupta, 2005; Strauss, personal communication, May 5, 2009; see also Roodenrys & Hinton, 2002).

Spoken word production: Facilitative effects of neighbors

Considering the opposite effects of lexical neighbors in visual and spoken word recognition, one might imagine that the critical difference is the modality itself; that something about spoken language makes neighbor effects inhibitory. This hypothesis is inconsistent with the robust facilitative effects of lexical neighbors on spoken word production (e.g., Gordon, 2002; Kittredge, Dell, Verkuilen, & Schwartz, 2008; Middleton & Schwartz, 2010; Vitevitch, 1997, 2002; Vitevitch & Sommers, 2003). These facilitative effects have been shown in error rates from natural speech corpora (e.g., Vitevitch, 1997), error elicitation tasks (e.g., Vitevitch, 2002), and picture naming tasks with healthy control participants (Mirman, Kittredge, & Dell, 2010; Vitevitch & Sommers, 2003) and aphasic speakers (e.g., Gordon, 2002; Kittredge et al., 2008; Middleton & Schwartz, 2010; Mirman et al., 2010). Similar patterns have also been observed in response times (e.g., Vitevitch & Sommers, 2003) and dense lexical neighborhoods seem to help in avoiding tip-of-the-tongue states (i.e., words from sparse neighborhoods are more likely to cause tip-of-the-tongue states; e.g., Harley & Brown, 1998; Vitevitch & Sommers, 2003).

As discussed above, phonological neighbors appear to exert opposite effects in visual and spoken word recognition. The same pattern does not appear to hold for word production: analyses of spoken and written spelling by aphasic participants showed that neighbors facilitated the successful production of a target word in either modality (Goldrick, Folk, & Rapp, 2010).

Facilitative effects of lexical neighbors in reading aloud could be due to facilitation at the visual word recognition level or the spoken word production level. However, researchers have argued that this effect is driven by the (phonological) word production aspect of the task rather than the (orthographic) word recognition aspect (Balota et al., 2004; Mulatti, Reynolds, & Besner, 2006).

Dell and Gordon (2003) described a preliminary computational account of why lexical neighbors facilitate spoken word production but inhibit spoken word recognition. Their account, which was implemented and tested in simulations of the two-step interactive-activation model of lexical access (Dell et al., 1997), was based on two key insights. The first was that phonological neighbor effects in spoken word production require interactivity; that is, bi-directional excitation between different levels, specifically between lexical and phonological representations. For phonologically similar facilitative “gangs” (cf. McClelland & Rumelhart, 1981; Taraban & McClelland, 1987) of words to become active, phoneme activation must feed back to lexical levels in order to activate them. The second was that word production is a semantically-driven task and word recognition is a phonologically-driven task. As a result, the strongest lexical competitors during word production are semantic neighbors, not phonological neighbors, and therefore the (weak) activation of some phonological neighbors should not substantially increase ambiguity. The notion that phonological neighbors facilitate word production by helping the target word overcome competition from semantic neighbors is also supported by reduced semantic and omission error proportions for high phonological neighborhood density words (Kittredge et al., 2008; see also Middleton & Schwartz, 2010). In contrast, phonological neighbors are the strongest competitors in spoken word recognition; therefore, having more of them (i.e., a dense neighborhood) should substantially increase ambiguity and slow down recognition. As we will show, this distinction between strong and weak competitors is at the core of determining whether neighbors will facilitate or inhibit processing.

Semantic neighbors: Opposite effects of near and distant neighbors

Following the same logic as defining neighbors in terms of form similarity (phonological or orthographic), neighbors can also be defined in terms of meaning (semantic) similarity. Several studies have found that words with many semantic neighbors or denser semantic neighborhoods are recognized more quickly (Buchanan, Westbury, & Burgess, 2001; Dunabeitia, Aviles, & Carreiras, 2008; Locker, Simpson, & Yates, 2003; Siakaluk, Buchanan, & Westbury, 2003; Yates, Locker, & Simpson, 2003). A finer-grain analysis of the effects of semantic neighbors on visual word recognition (concreteness judgment task) found that distant semantic neighbors – concepts that share a few semantic features – facilitated word recognition. In contrast, near semantic neighbors – concepts that share many semantic features – inhibited word recognition (Mirman & Magnuson, 2008). A subsequent study replicated this finding in a word production (picture naming) task testing aphasic and speeded control participants (Mirman, 2011).

Although these opposite effects of near and distant semantic neighbors have not been investigated as thoroughly as other neighbor effects, they provide important constraining evidence. Reversals across the other domains could possibly be attributed to differences in domain-specific representations (e.g., orthography vs. phonology) or tasks (e.g., word recognition vs. production). The semantic neighbor effects tell a different story: the neighbors are all of the same “type” (i.e., semantic) and their effects are consistent across tasks (both word recognition and production), but near semantic neighbors exert inhibitory effects while distant semantic neighbors exert facilitative effects.

The Present Study

Table 1 provides a summary of the qualitative effects of different kinds of lexical neighbors in different tasks, as reviewed above. Although greatly simplified, this summary captures the fact that neighborhood effects are considered among the most robust findings in each domain notwithstanding that, across domains, these robust effects go in opposite directions. This isolation of domains has arisen because most researchers have focused on only one kind of neighbor effect or context (e.g., spoken word recognition and not visual word recognition or spoken word production). However, they have almost universally appealed to interactive activation and competition in order to account for their findings, so a unified account may be possible. This is the goal of the present study: to empirically investigate the dynamics of interactive activation and competition in different word processing contexts in order to (a) examine whether the basic principles of interactive activation and competition (IAC) can correctly predict the direction of lexical neighbor effects and (b) uncover how or why the task and neighbor type may determine the direction of the effect.

Because our focus was specifically on the dynamics of interactive activation and competition, we used a very simple implementation of IAC principles in order to maximize the tractability of the simulations. Consequently, because of the minimal nature of the model and the importance of the qualitative reversals in the behavioral data, we focus on the qualitative patterns produced by the model (see also Pitt, Kim, Navarro, & Myung, 2006, for discussion of the value of global qualitative evaluation of computational models).

Simulations

Network architecture

Although the simple IAC model described here was not intended to be a full model of lexical processing, IAC is a core principle of PDP models in general, and as a result, the model architecture is closely related to leading models in each of the relevant domains.

Specifically, it is closely related to models of visual word recognition (e.g., Grainger & Jacobs, 1996; McClelland & Rumelhart, 1981), spoken word recognition (e.g., Gaskell & Marslen-Wilson, 1997; Magnuson, Tanenhaus, Aslin, & Dahan, 2003; McClelland & Elman, 1986), spoken word production (e.g., Dell et al., 1997), and lexical semantics (e.g., Cree, McRae, & McNorgan, 1999; Rogers & McClelland, 2004). Our modeling approach specifically builds on recent efforts to use simple IAC models to account for opposite effects of lexical neighbors on spoken word production vs. recognition (Dell & Gordon, 2003) and on spoken vs. visual word recognition (Magnuson & Mirman, 2007).

The basic structure of the model is shown in Figure 1. The model consisted of simple processing units organized into three layers: units in the first layer corresponded to elements of word form (i.e., phonemes or letters), units in the second layer corresponded to lexical elements (i.e., words in the model's lexicon), and units in the third layer corresponded to elements of meaning (i.e., semantic features of concepts denoted by the words). As in other IAC models (e.g., Dell et al., 1997; McClelland & Elman, 1986; McClelland & Rumelhart, 1981), congruent units in different layers were connected by bi-directional weighted connection links. That is, each word unit had bi-directional connections to its constituent letters or phonemes and to its semantic features.

To implement competition, units in the word layer were connected by bi-directional inhibitory connections. The inhibitory connection strength was scaled by a sigmoid function of unit activation, as shown in Figure 2. In other words, weakly active word units had very little inhibitory effect on other word units and strongly active words units had a very strong inhibitory effect on other word units. This non-linearity was implemented in order to allow initial parallel activation of many word candidates while still forcing the model to eventually settle to a single active representation. A similar approach was implemented by Cisek (2006) in a neurally-based computational model of action planning and selection. In essence, this approach implements a decision-making or response selection mechanism similar to the Luce (1959) choice rule (see also Mirman, Yee, Blumstein, & Magnuson, 2011, for a sigmoid representation of the Luce choice rule) and that captures the notion of progressively increasing pressure to settle to a single active candidate (for related discussion of progressively increasing decision pressure, see Allopenna, Magnuson, & Tanenhaus, 1998; McRae, Spivey, & Tanenhaus, 1998).

To mimic the kind of recurrent connections that would be learned by semantic representations (e.g., Cree et al., 1999; Rogers & McClelland, 2004), semantic units generally had inhibitory connections, but this inhibition was reduced for each concept in which the semantic units (features) co-occurred (for evidence of facilitative effects of feature co-occurrence, see e.g., Cree & McRae, 2003; Rogers & McClelland, 2004). In other words, a semantic feature such as “has wings” was assumed to have inhibitory connections to unrelated features such as “has strings”, excitatory connections to strongly (cor)related features such as “has feathers”, and intermediate weights to weakly related features such as “made of metal” (airplane). Implementation details of these within-level connections at the semantic level are described in Simulation 52. All units followed the standard IAC activation function in which positive net input drives unit activation toward its maximum (1.0) and negative net input drives unit activation toward its minimum (0.0). Complete model implementation details are provided in the Appendix, along with parameter values and ranges over which the simulated effects held. Full model code is available at: <https://sites.google.com/site/neighbormodel/model>.

²Simulations 1–4 were concerned with interactions between the form layer and the word layer during tasks in which the semantic properties of target words were controlled and balanced, so for simplicity, the semantic layer was omitted from these simulations.

All of the simulations used a simple lexicon consisting of five two-letter (or two-phoneme) words. Each word was also associated with 10 semantic feature units. Four of the words were neighbors (of whatever type was relevant for the simulation); the fifth word had no neighbors. The simulations compared processing of “word 1”, a high neighborhood density word, and “word 5”, the low neighborhood density word.

Simulations 1–3: Word recognition

As reviewed above, form neighbors tend to facilitate visual word recognition (e.g., Andrews, 1997; Yates, 2005), but have an inhibitory effect if they are higher in word frequency than the target word (Davis et al., 2009; Ferraro & Hansen, 2002; Grainger & Jacobs, 1996; Grainger et al., 1989, 1992; Grainger & Segui, 1990) or if the task is spoken word recognition (e.g., Luce & Pisoni, 1998; Magnuson et al., 2007). For these simulations, the words in the high density neighborhood shared their first letter or phoneme and the word in the low density neighborhood shared no letters or phonemes with any other word. To simulate processing of a printed word (Simulations 1 and 2), the units corresponding to its constituent letters were activated simultaneously and activation was allowed to propagate through the network until one of the word units crossed a response threshold value (0.7). Simulated response time (RT) was taken to be the number of time steps needed to reach this threshold. To simulate the effect of word frequency (Simulation 2), the connection weights between the higher frequency word units and the corresponding letter units were increased. This implementation was chosen based on evidence that word frequency affects the strength of the mappings represented by these connections (e.g., Dahan, Magnuson, & Tanenhaus, 2001) and because in models where such mappings are learned, the higher-frequency mappings are learned more strongly (e.g., Gaskell & Marslen-Wilson, 1997; Plaut et al., 1996). To simulate processing of a spoken word (Simulation 3) the same input was presented sequentially. The first phoneme was activated for time steps 1–30 and the second phoneme was activated for time steps 25–54. The small amount of overlap was intended as a rough analog to coarticulation. The simulation results did not depend on the precise amount of “coarticulatory” overlap.

In Simulation 1 (Figure 3, left), the visual word with many neighbors was recognized faster than the word with fewer neighbors. As shown in the figure, visual word neighbors were weakly and transiently activated. This weak activation was too low to cause substantial lexical inhibition (see Figure 2), but the bi-directional feedback to the form layer units provided additional excitation to the shared letter unit, which facilitated target recognition. This is precisely the “gang” or “conspiracy” effects described for other IAC models (e.g., McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982; Taraban & McClelland, 1987; see also Dell & Gordon, 2003) and proposed to account for neighborhood facilitation.

The results of Simulation 2 (Figure 3, middle) were also consistent with the behavioral data: the word with high frequency neighbors was processed more slowly than the word with equal frequency neighbors. Due to their stronger connections, higher frequency neighbors became active more quickly, reached a higher activation level, and remained active longer than equal frequency neighbors. Importantly, reaching a higher level of activation allowed them to inhibit the target word more strongly. Because inhibition strength was a nonlinear function of activation, this inhibitory effect out-weighed their increased recurrent facilitative effect on the shared letter unit. As the other simulations will show, this balance between the recurrent facilitative effects of neighbors and their lateral inhibitory effects – which hinges on their degree of activation – determines whether neighbors will exert a net facilitative or inhibitory effect on processing.

The change to serial input (Simulation 3) from parallel input (Simulation 1) reversed the net effect of lexical neighbors. When the input was presented serially, lexical neighbors exerted

an inhibitory effect on word processing (Figure 3, right). This comparison replicates the previously reported results from a very similar model that also showed opposite effects of neighbors for serial vs. parallel input (Magnuson & Mirman, 2007). Note that neighbors were more strongly activated by serial input than parallel input (compare diamonds in left and right panels of Figure 3). As in Simulation 2, when neighbors became more active, their inhibitory effect on the target word began to outweigh their facilitative recurrence with the form layer. In Simulations 2 and 3 the increased activation of neighbors was due to different causes – in Simulation 2 it was due to their higher frequency (stronger connections), in Simulation 3 it was due to lack of disambiguating information during the first phoneme (the target and the neighbors had the same first phoneme). Nevertheless, the outcome was the same: the facilitative effect of neighbors became inhibitory when the neighbors were activated more strongly.

Simulation 4: Word production

As reviewed above, lexical neighbors facilitate word production, making it faster and less error-prone (Goldrick et al., 2010; Gordon, 2002; Middleton & Schwartz, 2010; Mirman et al., 2010; Vitevitch, 1997, 2002; Vitevitch & Sommers, 2003). To simulate this task, the model from Simulation 3 was (conceptually) run in reverse: semantic input was presented to the word layer and allowed to propagate through the model; phoneme unit activations were treated as the model output (the same units could be considered grapheme output for simulating written word production, cf. Goldrick et al., 2010). Reaction time in word production studies is usually measured from stimulus onset to the onset of vocalization. To create a model analog of this measure, we considered the number of processing cycles from the onset of the semantic input until at least one phoneme (assumed to be the first phoneme³ in the vocalization) passed an activation threshold (0.7).

The results (Figure 4) were consistent with the behavioral data: word production was faster for a word with many neighbors than for a word with no neighbors. Feedback from the form layer partially activated the lexical units corresponding to the neighbors, but this activation was not enough to cause substantial inhibition of the target word. It was, however, enough to provide additional support to the shared phoneme in the target word, thus facilitating its production. One consequence of this is that the model predicts that neighbor effects should be largest on the neighbor-supported phonemes. This prediction is consistent with behavioral data showing that, even when the number of neighbors is matched, word production is faster when neighbors support more of the phoneme positions in the target word (Yates, Friend, & Ploetz, 2008). Although the effect of neighbors was largest on the shared phoneme, because of the recurrent feedback loop between phoneme and word units, the neighbor-supported phoneme unit also provided additional excitatory activation to the target word, which also facilitated activation of the other, not shared, phoneme. Simply put, facilitative neighbor effects are predicted to be strongest on the shared phonemes, but not limited to those phonemes.

The contrast between Simulations 3 and 4 (word recognition vs. word production) is consistent with simulations of a similar IAC model (Dell & Gordon, 2003). Dell and Gordon pointed out that phonological neighbors should be more strongly activated in word recognition, which is a phonologically-driven task, than in word production, which is a semantically-driven task. Our simulations are consistent with this claim and demonstrate

³This should not be taken to mean that the model's articulatory planning is strictly phoneme-by-phoneme. Both output phonemes were activated in parallel and the activation of the second phoneme lagged only slightly behind the first phoneme, reaching an average activation of 0.67 when the first phoneme reached the 0.7 threshold. This partial activation is broadly consistent with the observation that in natural speech there is substantial coarticulation between phonemes.

that this degree of activation difference determines whether the net effect of neighbors will be facilitative (weakly active neighbors) or inhibitory (strongly active neighbors).

Simulations 5 and 6: Effects of near and distant semantic neighbors

Simulations 1–4 focused on effects of form neighbors and on reversals in those effects due to task or stimulus differences. In the domain of semantic neighbors, such reversals have been demonstrated within task. Specifically, both visual word recognition (Mirman & Magnuson, 2008) and word production (Mirman, 2011) were slower for words with many near semantic neighbors (concepts that share many semantic features) and faster for words with many distant semantic neighbors (concepts that share a few features). Since the materials in those behavioral experiments were matched on orthographic and phonological neighborhood, for simplicity of simulations, only the word and semantic layers were included in these simulations. To simulate word recognition (Simulation 5), input was presented directly to the word layer and allowed to propagate bi-directionally between the word layer and the semantic layer. To simulate word production (Simulation 6), input was presented to each of the semantic feature units corresponding to the target word.

Each word unit was connected to 10 semantic feature units. Lateral connections among semantic feature units were set on the basis of whether those features tended to co-occur across concepts. units that never occurred together in a concept were connected by negative weights (-0.03), feature units that sometimes occurred together and sometimes separately were connected by small positive weights (0.002), and feature units that always occurred together were connected by positive weights whose magnitude was a function of the number of co-occurrences (C ; e.g., Cree & McRae, 2003; Rogers & McClelland, 2004):

$$W(C)=0.016+0.004*C$$

Near semantic neighbors were defined as sharing 8 out of 10 semantic features; distant semantic neighbors were defined as sharing 4 out of 10 semantic features. In the behavioral studies (Mirman, 2011; Mirman & Magnuson, 2008) the term "many" meant a quite different number for near and distant neighbors: for near neighbors it was around 4, but for distant neighbors it was over 200. To capture this distinction, the word with "many" near neighbors had 1 near neighbor and for the word with "many" distant neighbors had 10 distant neighbors. As in the previous simulations, the word with "few" neighbors had 0 neighbors (near or distant). For Simulation 5, because the behavioral study used a semantic task (concreteness judgment), we considered the semantic layer to be the "output" layer. However, because of the distributed semantic representation, there was no single unit whose activation would correspond to target (or neighbor) activation. As a result, word recognition was assessed using normalized cross-entropy error (nCEE; Mirman & Magnuson, 2008; see Appendix for more details) to measure the distance in semantic state space between the model's state and ideal states corresponding to target and neighbor words. Reaction time was measured as the number of time steps from stimulus onset (start of input to word layer) until the nCEE dropped below a threshold (0.2), i.e., the number of time steps required for the activation pattern in the semantic layer to get close enough to the target word to consider the word to have been recognized. For Simulation 6 (word production), the model's task performance was based directly on activations of the word layer units and response time computed as the number of time steps required for the target word unit to reach the response threshold (0.7 , as in the other simulations).

The simulation results (Figure 5) showed that the model exhibited opposite effects of near and distant semantic neighbors on word recognition (Simulation 5, left panels in Figure 5) and word production (Simulation 6, right panels in Figure 5). In Simulation 5, the model

was slower to settle to the target representation when the target word had many near semantic neighbors than when it had none (Figure 5: top left panel) and faster when the target word had many distant semantic neighbors than when it had none (Figure 5: bottom left panel). In Simulation 6, word activation was slower for words with many near semantic neighbors (Figure 5: top right panel) and faster for words with many distant semantic neighbors (Figure 5: bottom right panel). As in the other simulations, strongly activated neighbors (near neighbors) had a net inhibitory effect on word recognition and weakly activated neighbors (distant neighbors) had a net facilitative effect on word recognition.

General Discussion

Summary of Key Findings

One of the most widely agreed-upon principles in cognition is that multiple similar representations are activated in parallel and compete for selection. This principle plays a key role in theories and models across a diverse set of domains, including perception (e.g., Palmeri, Wong, & Gauthier, 2004), action (e.g., Botvinick, Buxbaum, Bylsma, & Jax, 2009; Cisek, 2006), categorization (e.g., Kalish et al., 2004), memory (e.g., Polyn, Norman, & Kahana, 2009), and language (e.g., McClelland & Rumelhart, 1981). In the language domain this principle has been studied extensively in the context of “neighborhood” effects – how recognition or production of a target word is affected by words that are similar to it. Orthographic, phonological, and semantic similarity have been considered in a wide variety of tasks (picture naming, word reading, word repetition, lexical decision, semantic categorization or judgment, etc.). Across these many studies a striking pattern of consistent reversals has emerged: given a particular task and neighbor type, the effects are quite consistent, but the direction of the effect – facilitation vs. inhibition – differs across tasks and neighbor types. Although neighbor effects are one of the most robust findings in lexical processing tasks, there has been little effort to explain why the same neighbors would, for example, have facilitative effects on spoken word production, inhibitory effects on spoken word recognition, and facilitative effects on visual word recognition. Further, the accounts of individual facilitative or inhibitory effects have almost universally appealed to interactive activation and competition, without addressing why the same framework would predict opposite effects in these different contexts.

Here, we have addressed this specific question using a simple implementation of the IAC framework. We deliberately chose a simple version of IAC that did not strive to capture all of the details of any one task so that we could use the same model across tasks. The simulation results captured the core qualitative patterns of orthographic, phonological, and semantic neighbor effects in word recognition and production tasks (summarized in Table 1). By using the same model with the same parameter values across simulations, we were able to extract a core computational principle that determined whether neighbor effects were facilitative or inhibitory: *strongly active neighbors had a net inhibitory effect and weakly active neighbors had a net facilitative effect*. This pattern emerged from comparisons across simulations and was not dependent on specific parameter settings (i.e., the qualitative simulation results held over fairly large changes in parameter values, see Appendix for details). This pattern is also consistent with the few previous attempts to explain why neighbors have opposite effects in different tasks (Dell & Gordon, 2003; Magnuson & Mirman, 2007).

The qualitative results demonstrated in the present simulations critically depend on the sigmoid inhibition function implemented at the word layer. This implementation was rooted in standard models of decision processes, most notably the Luce (1959) choice rule (see also Usher & McClelland, 2001), and in neural evidence (e.g., Cisek, 2006). Further, in Simulation 4 (effects of phonological neighbors on word production) and Simulation 5

(effects of semantic neighbors on word recognition), the model output was not read directly from the word layer, so idiosyncratic effects limited to the word layer could not explain those results. In other words, the sigmoid inhibition function played a critical role in the model dynamics, not just word level dynamics. Finally, when considered from the perspective of explanatory power, evidence that this single (well-motivated) processing principle can account for a large and complex set of qualitative data patterns, and do so over a relatively large parameter range, suggests that this principle may be an important aspect of cognitive processing. On this view, the simulation results reported here serve as an existence proof demonstrating that a diverse set of findings can be explained by an underlying nonlinear (sigmoid) relationship between activation and inhibition of competing representations.

Limitations, Speculations, and Future Directions

Effects of lexical neighbors on word processing are among the most studied phenomena in lexical processing and highly detailed models have been developed to account for individual kinds of neighborhood effects, such as orthographic neighbors in visual word recognition. However, so far, there has been little effort to develop a unified account across tasks and neighbor types. Here we have taken a step toward developing such an account. Because the goal was to examine domain-general computational properties, we used a very simple and general model that was designed to capture the principles of interactive activation and competition in a way that is consistent with existing models of written and spoken language processing (Dell et al., 1997; Jacobs & Grainger, 1994; McClelland & Elman, 1986; McClelland & Rumelhart, 1981), and general enough to be applied to very different tasks. That is, our model was meant to be the beginning of a bridge across domain-specific accounts of neighborhood effects.

Domain-specific models have provided very detailed (even item-level) accounts of neighborhood effects in their domains, but such models are limited in that they do not speak to any neighbor effects in other domains. Our model addresses effects of neighbors across domains (indeed, the model was general enough that it is possible to re-label it to be a model of object recognition, memory, categorization, or any other domain where interactive activation and competition are proposed as core mechanisms), but it is limited in that it does not capture important domain-specific factors. We believe our simulations provide important complementary evidence that will help to guide the integration of domain-specific models into a domain-general account that captures both the strengths of our simple domain-general model and existing domain-specific models.

A second, related limitation is that the present simulations used very small lexicons with very simple neighbor relations. As a result, the domain-generality that we observed in our simulations may have been bolstered by the similarity of the neighborhood structures that we used. Studies with larger-scale models that build on our model and implement more domain-specific details are needed to investigate the balance between domain-specific vs. domain-general properties of neighbor effects.

Although our model used a simple neighbor definition, the general principle that it demonstrates is applicable to effects of neighbors defined in other ways. One such case is the inhibitory effect of transposed-letter neighbors in visual word recognition (Andrews, 1996): words like *salt*, which has the transposed-letter neighbor *slat*, are recognized more slowly than matched words that do not have a transposed-letter neighbor (such as *halt*). This pattern could arise if transposed-letter neighbors are more confusable than substituted-letter neighbors, which have facilitative effects on visual word recognition; that is, if transposed-letter neighbors are stronger competitors than substituted-letter neighbors. This would be analogous to the results of Simulations 5 and 6, which showed inhibitory effects of near

semantic neighbors and facilitative effects of distant semantic neighbors, and to Simulation 2, which showed that high-frequency neighbors inhibit target word recognition because they are more active.

Another important consideration is the clustering or spread of neighbors, which has been shown to affect spoken (Chan & Vitevitch, 2009; Vitevitch, 2007) and visual (Mathey & Zagar, 2000) word recognition and word production (Yates et al., 2008). We have focused on the effects of neighbors on target word processing, but the neighbors also affect one another, which should have indirect but measurable effects on target processing. Our simulations predict that if particular patterns of neighbor clustering lead the neighbors to enhance one another's activation, then they will tend to have more inhibitory (or less facilitative) effects; in contrast, if they do not accentuate one another's activation, then their cumulative effect on the target will be more facilitative. This could explain why clustered neighbors (i.e., neighbors that are also neighbors of one another) are particularly inhibitory in spoken word recognition (Chan & Vitevitch, 2009).

A related question is whether phonotactic probability facilitates spoken word recognition, in contrast with the inhibitory effects of lexical neighbors in the same conditions (Luce & Large, 2001; Vitevitch & Luce, 1998, 1999; Vitevitch et al., 1999). Our simulations suggest that this pattern could arise if phonotactic probability is indirectly measuring "distant" phonological neighborhoods; that is, if high phonotactic probability words cause diffuse, weak activation of lexical representations, which would have a net facilitative effect. This would contrast with traditional lexical neighborhood metrics (such as the 1-phoneme rule), which capture "near" phonological neighbors that are strongly activated and would have a net inhibitory effect. Perhaps the difficulty in replicating the phonotactic probability effect (e.g., Lipinski & Gupta, 2005; Roodenrys & Hinton, 2002; Strauss, personal communication, May 5, 2009) stems from it indirectly measuring distant phonological neighbors. A more direct manipulation may provide clearer insights into this issue.

The present simulations relied on a simple "hard-wired" model that had no learning process and no intrinsic noise during processing. Thus, the model did not address how neighborhood effects interact with learning, development, neurological impairment, and cognitive decline due to normal aging. Similarly, the model did not address errors in word recognition and production. These are interesting and important issues for future research.

Finally, the present simulations relied on qualitative, rather than quantitative, comparisons of model and behavioral data. We chose this approach because there were such striking and unexplained qualitative reversals in the literature. By choosing a simple model, we were able to use the same model to qualitatively account for the full set of behavioral data and to extract global (relatively parameter-independent) insights into the dynamics of interactive activation and competition. It would be a mistake to expect such a simple model, chosen for domain-general rather than item-level specificity, to provide quantitative fits to behavioral data (see Pitt et al., 2006, for a discussion of different model evaluation methods). Rather, the basic principles explored here can form the basis for elaborated models that can, and should, be evaluated quantitatively.

Concluding Remarks

We have presented a series of simulations exploring whether the dynamics of interactive activation and competition can account for facilitative and inhibitory effects lexical neighbors. We chose to explore this issue in the domain of single word processing, but the computational principle is one of the most general in theories of cognition: parallel activation of multiple similar representations and competition (selection) among them. Single word processing provided an ideal context for investigating this issue because there is

a large literature showing a complex pattern of facilitative and inhibitory effects, which have been robustly replicated and explained individually, but not together. We used a simple interactive activation and competition model that could be applied to the full range of behavioral data without changing the model architecture or parameters. The simulations showed that the complex pattern of contrasting neighbor effects boils down to a simple computational principle: co-activated representations have both facilitative and inhibitory effects; they have a net inhibitory effect if they are strongly activated and a net facilitative effect if they are weakly activated. Because the model is so general, this insight applies to any domain where interactive activation and competition are thought to be involved; that is, any mental activity involving the processing of related or similar representations in a multi-level system – which covers much of perception, cognition, and action.

Acknowledgments

This research was supported by R01DC010805 to DM and by the Moss Rehabilitation Research Institute. The authors thank Gary Dell for helpful comments on an early draft.

References

- Acha J, Perea M. The effect of neighborhood frequency in reading: Evidence with transposed-letter neighbors. *Cognition*. 2008; 108:290–300. [PubMed: 18374323]
- Alloppenna PD, Magnuson JS, Tanenhaus MK. Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory & Language*. 1998; 38(4):419–439.
- Andrews S. Frequency and neighborhood effects on lexical access: Activation or search? *Journal of Experimental Psychology: Learning, Memory, and Cognition*(5). 1989; 15:802–814.
- Andrews S. Frequency and neighborhood effects on lexical access: Lexical similarity or orthographic redundancy? *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1992; 18:234–254.
- Andrews S. Lexical retrieval and selection processes: Effects of transposed-letter confusability. *Journal of Memory and Language*. 1996; 35(6):775–800.
- Andrews S. The effect of orthographic similarity on lexical retrieval: Resolving neighborhood conflicts. *Psychonomic Bulletin & Review*. 1997; 4:439–461.
- Balota DA, Cortese MJ, Sergent-Marshall SD, Spieler DH, Yap M. Visual Word Recognition of Single-Syllable Words. *Journal of Experimental Psychology: General*. 2004; 133(2):283–316. [PubMed: 15149254]
- Benki J. Quantitative evaluation of lexical status, word frequency, and neighborhood density as context effects in spoken word recognition. *Journal of the Acoustical Society of America*. 2003; 113:1689–1705. [PubMed: 12656401]
- Botvinick MM, Buxbaum LJ, Bylsma LM, Jax SA. Toward an integrated account of object and action selection: A computational analysis and empirical findings from reaching-to-grasp and tool-use. *Neuropsychologia*. 2009; 47(3):671–683. [PubMed: 19100758]
- Boyczuk JP, Baum SR. The influence of neighborhood density on phonetic categorization in aphasia. *Brain and Language*. 1999; 67(1):46–70. [PubMed: 10191000]
- Buchanan L, Westbury C, Burgess C. Characterizing semantic space: Neighborhood effects in word recognition. *Psychonomic Bulletin & Review*. 2001; 8(3):531–544. [PubMed: 11700905]
- Chan KY, Vitevitch MS. The influence of the phonological neighborhood clustering coefficient on spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*. 2009; 35(6):1934–1949. [PubMed: 19968444]
- Cisek P. Integrated neural processes for defining potential actions and deciding between them: A computational model. *The Journal of Neuroscience*. 2006; 26(38):9761–9770. [PubMed: 16988047]

- Cohen JD, Dunbar K, McClelland JL. On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*. 1990; 97(3):332–361. [PubMed: 2200075]
- Coltheart, M.; Davelaar, E.; Jonasson, J.; Besner, D. Access to the internal lexicon. In: Dornic, S., editor. *Attention and performance VI*. Hillsdale, NJ: Erlbaum; 1977. p. 535-555.
- Cree GS, McRae K. Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*. 2003; 132(2):163–201. [PubMed: 12825636]
- Cree GS, McRae K, McNorgan C. An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*. 1999; 23(3):371–414.
- Dahan D, Magnuson JS, Tanenhaus Michael K. Time course of frequency effects in spoken-word recognition: Evidence from eye movements. *Cognitive Psychology*. 2001; 42(4):317–367. [PubMed: 11368527]
- Davis CJ, Perea M, Acha J. Re(de)fining the orthographic neighborhood: the role of addition and deletion neighbors in lexical decision and reading. *Journal of Experimental Psychology: Human Perception and Performance*. 2009; 35(5):1550–1570. [PubMed: 19803656]
- Dell, GS.; Gordon, JK. Neighbors in the Lexicon: Friends or foes?. In: Schiller, NO.; Meyer, AS., editors. *Phonetics and phonology in language comprehension and production: Differences & similarities*. Mouton de Gruyter; 2003. p. 9-37.
- Dell GS, Schwartz MF, Martin N, Saffran EM, Gagnon DA. Lexical access in aphasic and nonaphasic speakers. *Psychological Review*. 1997; 104(4):801–838. [PubMed: 9337631]
- Dufour S, Peereman R. Inhibitory priming effects in auditory word recognition: When the target's competitors conflict with the prime word. *Cognition*. 2003a; 88(3):B33–B44. [PubMed: 12804820]
- Dufour S, Peereman R. Lexical competition in phonological priming: Assessing the role of phonological match and mismatch lengths between primes and targets. *Memory & Cognition*. 2003b; 31(8):1271–1283.
- Duñabeitia JA, Avilés A, Carreiras M. NoA's Ark: influence of the number of associates in visual word recognition. *Psychonomic Bulletin & Review*. 2008; 15(6):1072–1077. [PubMed: 19001569]
- Ferraro FR, Hansen CL. Orthographic neighborhood size, number of word meanings, and number of higher frequency neighbors. *Brain and Language*. 2002; 82(2):200–205. [PubMed: 12096876]
- Forster KI, Shen D. No enemies in the neighborhood: Absence of inhibitory neighborhood effects in lexical decision and semantic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1996; 22(3):696–713.
- Garlock VM, Walley AC, Metsala JL. Age-of-Acquisition, Word Frequency, and Neighborhood Density Effects on Spoken Word Recognition by Children and Adults. *Journal of Memory and Language*. 2001; 45(3):468–492.
- Gaskell MG, Marslen-Wilson WD. Integrating form and meaning: A distributed model of speech perception. *Language and Cognitive Processes*. 1997; 12(5/6):613–656.
- Goldinger SD, Luce PA, Pisoni DB. Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory & Language*. 1989; 28:501–518.
- Goldrick M, Folk JR, Rapp B. Mrs. Malaprop's Neighborhood: Using Word Errors to Reveal Neighborhood Structure. *Journal of Memory and Language*. 2010; 62(2):113–134. [PubMed: 20161591]
- Gordon JK. Phonological neighborhood effects in aphasic speech errors: Spontaneous and structured contexts. *Brain and Language*. 2002; 82(2):113–145. [PubMed: 12096871]
- Grainger J. Cracking the orthographic code: An introduction. *Language and Cognitive Processes*. 2008; 23(1):1–35.
- Grainger J, Jacobs AM. Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*. 1996; 103(3):518–565. [PubMed: 8759046]
- Grainger J, O'Regan JK, Jacobs AM, Segui J. On the role of competing word units in visual word recognition: the neighborhood frequency effect. *Perception & Psychophysics*. 1989; 45(3):189–195. [PubMed: 2710616]

- Grainger J, O'Regan JK, Jacobs AM, Segui J. Neighborhood frequency effects and letter visibility in visual word recognition. *Perception & Psychophysics*. 1992; 51(1):49–56. [PubMed: 1549424]
- Grainger J, Segui J. Neighborhood frequency effects in visual word recognition: a comparison of lexical decision and masked identification latencies. *Perception & Psychophysics*. 1990; 47(2): 191–198. [PubMed: 2304817]
- Griffiths, TL.; Kemp, C.; Tenenbaum, JB. Bayesian models of cognition. In: Sun, R., editor. *Cambridge Handbook of Computational Psychology*. Cambridge University Press; 2008. p. 59-100.
- Harley TA, Brown HE. What causes a tip-of-the-tongue state? Evidence for lexical neighbourhood effects in speech production. *British Journal of Psychology*. 1998; 89(1):157–174.
- Johnson RL. The quiet clam is quite calm: Transposed-letter neighborhood effects on eye movements during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2009; 35(4):943–969.
- Johnson NF, Pugh KR. A cohort model of visual word recognition. *Cognitive Psychology*. 1994; 26(3):240–346. [PubMed: 8033536]
- Kalish ML, Lewandowsky S, Kruschke JK. Population of Linear Experts: Knowledge Partitioning and Function Learning. *Psychological Review*. 2004; 111(4):1072–1099. [PubMed: 15482074]
- Kittredge AK, Dell GS, Verkuilen J, Schwartz MF. Where is the effect of frequency in word production? Insights from aphasic picture-naming errors. *Cognitive Neuropsychology*. 2008; 25(4):463–492. [PubMed: 18704797]
- Lipinski J, Gupta P. Does neighborhood density influence repetition latency for nonwords? Separating the effects of density and duration. *Journal of Memory and Language*. 2005; 52(2):171–192.
- Locker L, Simpson GB, Yates M. Semantic neighborhood effects on the recognition of ambiguous words. *Memory & cognition*. 2003; 31(4):505–515.
- Luce, PA. *Research on Speech Perception*, Technical Report No. 6. Bloomington, IN: Speech Research Laboratory, Department of Psychology, Indiana University; 1986. Neighborhoods of words in the mental lexicon.
- Luce PA, Large NR. Phonotactics, density, and entropy in spoken word recognition. *Language and Cognitive Processes*. 2001; 16(5/6):565–581.
- Luce PA, Pisoni DB. Recognizing spoken words: The neighborhood activation model. *Ear and Hearing*. 1998; 19(1):1–36. [PubMed: 9504270]
- Luce, RD. *Individual choice behavior*. Oxford, England: John Wiley; 1959.
- Magnuson JS, Dixon JA, Tanenhaus MK, Aslin RN. The dynamics of lexical competition during spoken word recognition. *Cognitive Science*. 2007; 31:1–24. [PubMed: 21635286]
- Magnuson, JS.; Mirman, D. Neighborhood effects in word recognition: It's not where you live, it's how you get home; Paper presented at the 48th Annual Meeting of the Psychonomic Society; Long Beach, CA. 2007.
- Magnuson JS, Tanenhaus MK, Aslin RN, Dahan D. The time course of spoken word learning and recognition: Studies with artificial lexicons. *Journal of Experimental Psychology: General*. 2003; 132(2):202–227. [PubMed: 12825637]
- Mathey S, Zagar D. The neighborhood distribution effect in visual word recognition: words with single and twin neighbors. *Journal of Experimental Psychology: Human Perception and Performance*. 2000; 26(1):184–205. [PubMed: 10696613]
- McClelland, JL. Toward a Theory of Information Processing in Graded, Random, and Interactive Networks. In: Meyer, DE.; Kornblum, S., editors. *Attention & Performance XIV: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience*. Cambridge: MIT Press; 1993. p. 655-688.
- McClelland JL, Elman JL. The TRACE model of speech perception. *Cognitive Psychology*. 1986; 18(1):1–86. [PubMed: 3753912]
- McClelland JL, Rumelhart DE. An interactive activation model of context effects in letter perception: I An account of basic findings. *Psychological Review*. 1981; 88(5):375–407.
- McRae K, Spivey-Knowlton MJ, Tanenhaus Michael K. Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*. 1998; 38(3):283–312.

- Medin DL, Schaffer MM. Context theory of classification learning. *Psychological Review*. 1978; 85:207–238.
- Middleton EL, Schwartz MF. Density pervades: An analysis of phonological neighbourhood density effects in aphasic speakers with different types of naming impairment. *Cognitive neuropsychology*. 2010; 27(5):401–427. [PubMed: 21718214]
- Mirman D. Effects of near and distant semantic neighbors on word production. *Cognitive, Affective, and Behavioral Neuroscience*. 2011; 11(1):32–43.
- Mirman D, Dixon JA, Magnuson JS. Statistical and computational models of the visual world paradigm: Growth curves and individual differences. *Journal of Memory and Language*. 2008; 59(4):475–494. [PubMed: 19060958]
- Mirman, D.; Kittredge, AK.; Dell, GS. Effects of Near and Distant Phonological Neighbors on Picture Naming; Proceedings of the 32nd Annual Cognitive Science Society Meeting; 2010. p. 1447-1452.
- Mirman D, Magnuson JS. Attractor dynamics and semantic neighborhood density: Processing is slowed by near neighbors and speeded by distant neighbors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2008; 34(1):65–79.
- Mirman D, Yee E, Blumstein SE, Magnuson JS. Theories of spoken word recognition deficits in aphasia: Evidence from eye-tracking and computational modeling. *Brain and Language*. 2011; 117(2):53–68. [PubMed: 21371743]
- Mulatti C, Reynolds MG, Besner D. Neighborhood effects in reading aloud: new findings and new challenges for computational models. *Journal of Experimental Psychology: Human Perception and Performance*. 2006; 32(4):799–810. [PubMed: 16846280]
- Newman RS, Sawusch JR, Luce PA. Lexical neighborhood effects in phonetic processing. *Journal of Experimental Psychology: Human Perception & Performance*. 1997; 23(3):873–889. [PubMed: 9180048]
- Palmeri TJ, Wong AC-N, Gauthier I. Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences*. 2004; 8(8):378–386. [PubMed: 15335465]
- Pitt MA, Kim W, Navarro DJ, Myung JI. Global model analysis by parameter space partitioning. *Psychological Review*. 2006; 113(1):57–83. [PubMed: 16478301]
- Plaut DC, McClelland JL, Seidenberg MS, Patterson K. Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*. 1996; 103(1): 56–115. [PubMed: 8650300]
- Polyn SM, Norman KA, Kahana MJ. A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*. 2009; 116(1):129–156. [PubMed: 19159151]
- Rogers, TT.; McClelland, JL. *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: MIT Press; 2004.
- Roodenrys S, Hinton M. Sublexical or lexical effects on serial recall of nonwords? *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2002; 28(1):29–33.
- Rumelhart DE, McClelland JL. An interactive activation model of context effects in letter perception: II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*. 1982; 89(1):60–94. [PubMed: 7058229]
- Rumelhart, DE.; McClelland, JL. the PDP Research Group. *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press; 1986.
- Sears CR, Hino Y, Lupker SJ. Neighborhood size and neighborhood frequency effects in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*. 1995; 21(4):876–900.
- Siakaluk PD, Buchanan L, Westbury C. The effect of semantic distance in yes/no and go/no-go semantic categorization tasks. *Memory & Cognition*. 2003; 31(1):100–113.
- Taraban R, McClelland JL. Conspiracy effects in word pronunciation. *Journal of Memory and Language*. 1987; 26(6):608–631.
- Usher M, McClelland JL. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*. 2001; 108(3):550–592. [PubMed: 11488378]
- Vitevitch MS. The neighborhood characteristics of malopropisms. *Language and Speech*. 1997; 40(3): 211–228. [PubMed: 9509578]

- Vitevitch MS. The influence of phonological similarity neighborhoods on speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2002; 28(4):735–747.
- Vitevitch MS. The spread of the phonological neighborhood influences spoken word recognition. *Memory & Cognition*. 2007; 35(1):166–175.
- Vitevitch MS, Luce PA. When words compete: Levels of processing in perception of spoken words. *Psychological Science*. 1998; 9(4):325–329.
- Vitevitch MS, Luce PA. Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory & Language*. 1999; 40(3):374–408.
- Vitevitch MS, Luce PA, Pisoni DB, Auer ET. Phonotactics, neighborhood activation, and lexical access for spoken words. *Brain and Language*. 1999; 68(1–2):306–311. [PubMed: 10433774]
- Vitevitch MS, Sommers MS. The facilitative influence of phonological similarity and neighborhood frequency in speech production in younger and older adults. *Memory & Cognition*. 2003; 31(4): 491–504.
- Yarkoni T, Balota D, Yap M. Moving beyond Coltheart’s N: A new measure of orthographic similarity. *Psychonomic Bulletin & Review*. 2008; 15(5):971–979. [PubMed: 18926991]
- Yates M. Phonological Neighbors Speed Visual Word Processing: Evidence From Multiple Tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2005; 31(6):1385–1397.
- Yates M, Friend J, Ploetz DM. Phonological neighbors influence word naming through the least supported phoneme. *Journal of Experimental Psychology: Human Perception and Performance*. 2008; 34(6):1599–1608. [PubMed: 19045995]
- Yates M, Locker L, Simpson GB. Semantic and phonological influences on the processing of words and pseudohomophones. *Memory & Cognition*. 2003; 31(6):856–866.
- Yates M, Locker L, Simpson GB. The influence of phonological neighborhood on visual word perception. *Psychonomic Bulletin & Review*. 2004; 11(3):452–457. [PubMed: 15376794]

Appendix

The model was a simple “neural network” implementation of the interactive activation and competition framework (e.g., McClelland & Rumelhart, 1981). The network was composed of simple processing units organized into layers. Each unit had an activation value and was connected to other units by weighted connections, which could be excitatory or inhibitory. At the start of each simulation, all units were initialized to the rest activation values (0) and external input was provided to the appropriate units. This activation was then propagated by the weighted connections over a series of time steps. On each time step, each unit computed its net input from all units connected to it on the basis of their activations. Specifically, the net input to unit i was:

$$net_i = extinput_i + \sum_j w_{ij} a_j$$

where w_{ij} is the connection weight from unit j to unit i , a_j is the activation of unit j , and $extinput_i$ is any external input to unit i . Every unit’s activation was then updated as follows:

$$\begin{aligned} \text{If } (net_i > 0), \Delta a_i &= (max - a_i) net_i - decay(a_i - rest) \\ \text{Otherwise, } \Delta a_i &= (a_i - min) net_i - decay(a_i - rest) \end{aligned}$$

where max is the maximum activation, min is the minimum activation, $rest$ is the resting activation level, and $decay$ is a constant that brings the activation of the unit back to resting activation level. In all simulations reported in this paper, we choose $max=1$, $min=0$, and $rest=0$. These are typical values used in many interactive activation and competition network models. See Table A1 for the full list of other model parameters and their values.

At the word layer, the inhibitory weights were scaled by a sigmoid function of their activation:

$$y = \frac{15}{1.5 + e^{-\beta(x-x_0)}}$$

where increasing the value of the parameter β increases the steepness of the curve, and x_0 determines the “cross-over” point (i.e., the activation value at which the inhibition weight scaling factor is half-way between its minimum and maximum). For all simulations reported here, $\beta = 35$ and $x_0 = 0.3$; the sigmoid curve corresponding to these parameter values is shown in Figure 2. This sigmoid scaling allows multiple word units to be activated initially (inhibition is weak when unit activation is low) and forces the model to rapidly settle to a single active word unit (rising activation causes a fast increase in inhibition strength; see also, Cisek, 2006).

The qualitative simulation results were quite robust over a large range of parameter values. The results required a sigmoid scaling of the inhibition strength between words, but the particular parameters of the sigmoid function were not critical. For the parameters in Table A1, almost an entire $\pm 50\%$ range of the given values produced all of the same qualitative patterns. One relatively sensitive parameter was the frequency scaling factor in Simulation 2 (increased by 36%), which determined the strength of letter-to-word excitation weights for higher frequency words relative to lower frequency words. This parameter needed to be balanced with the model’s other excitatory and inhibitory weights in order to achieve the correct pattern of recurrent facilitative effects and lateral inhibitory effects. Note that this was a matter of balance, not absolute value: for any choice of the parameters in Table A1, there was a range for this scaling factor that produced the reported result, but this range was specific to those other parameter values.

Table A1

Parameter values used in the all simulations

Parameter	Value
phoneme/letter to word excitation	0.1
word to phoneme/letter excitation	0.1
word to semantics excitation	0.03
semantics to word excitation	0.03
word to word inhibition	0.04
phoneme/letter decay	0.01
word layer decay	0.01
semantic layer decay	0.05

Note. The excitation and inhibition parameters refer to connection weights.

Cross-entropy calculation

In Simulation 5, the output was read from the distributed semantic representation, which could not be evaluated based on activation of a single unit. We used cross-entropy error (CEE) to assess the distance between the activation pattern in the semantic layer and the target activation pattern. The CEE function is:

$$CEE = - \sum_j d_j \ln(s_j) + (1 - d_j) \ln(1 - s_j)$$

where d_j is the target activation for the j th semantics unit, and s_j is the current activation for that unit. This raw CEE value was normalized (divided by the maximum CEE for each item), to remove the effects of model starting state and map the values into the same 0–1 scale as the other simulations. Conceptually, normalized CEE represents the proportion of the distance between the model's starting state and the target state that the model has traversed.




Figure 1.
The full structure for the model

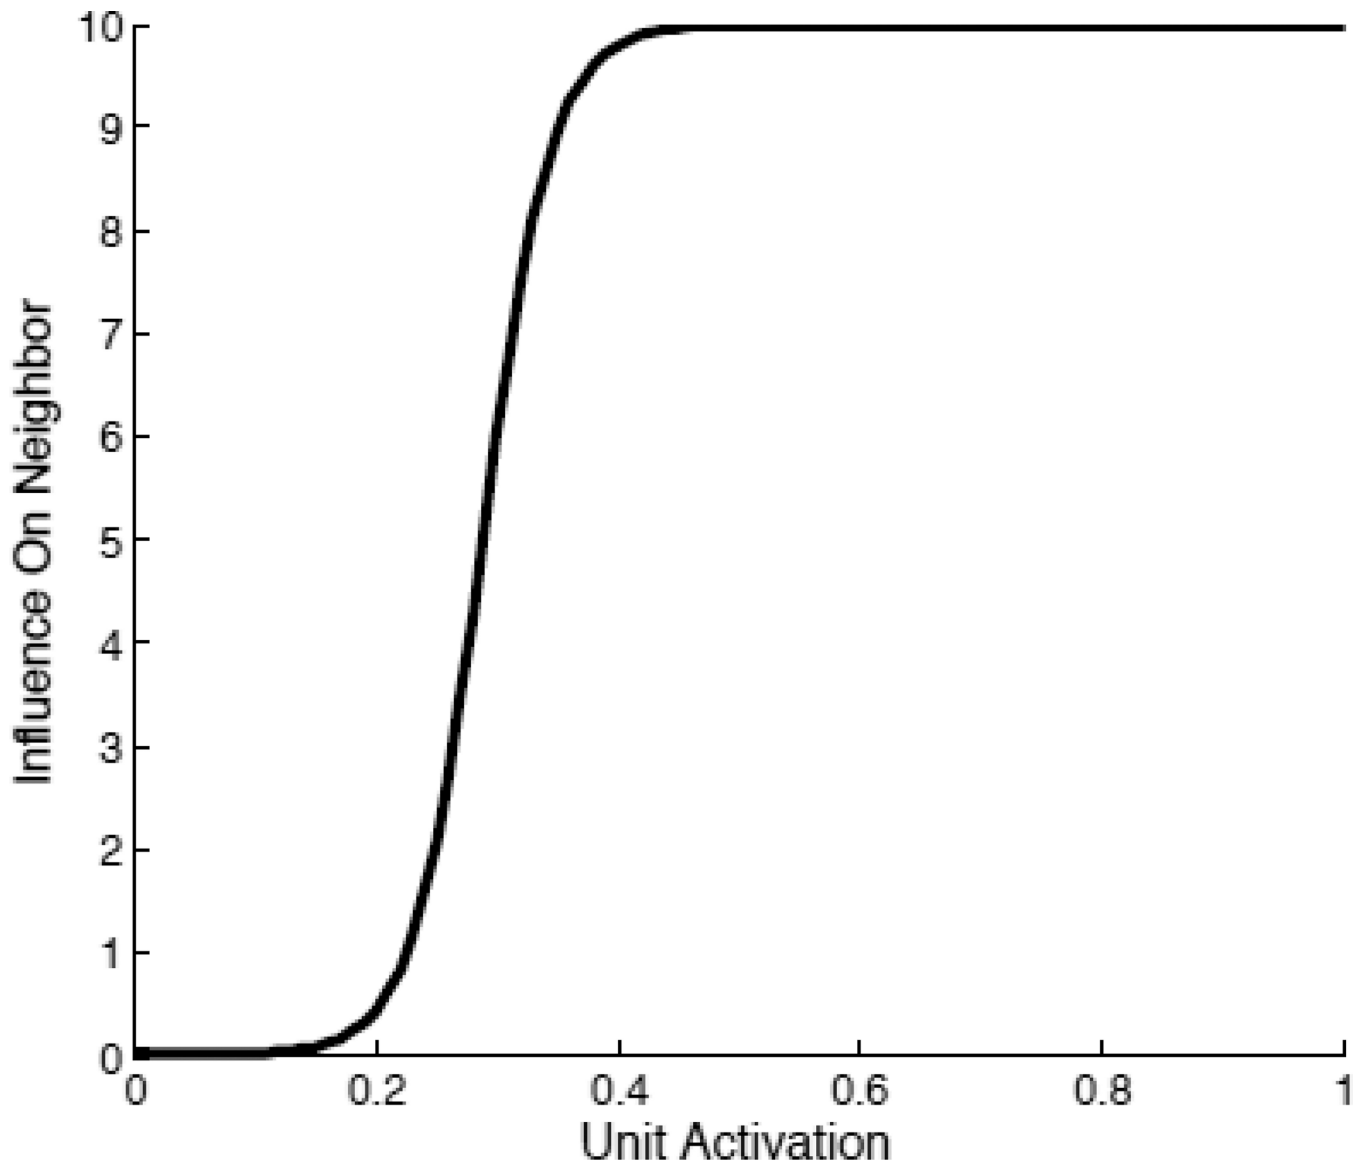


Figure 2.
Sigmoid inhibition strength function




Figure 3.

Results of Simulations 1–3. The left panel shows Simulation 1 (visual word recognition), the middle panel shows Simulation 2 (neighbor frequency effect on visual word recognition), and the right panel shows Simulation 3 (spoken word recognition). The curves represent the time course of activation for high neighborhood density target words (crosses), low neighborhood density target words (squares), target words with higher frequency neighbors (open circles). Activation of the equal frequency neighbors (filled circles) and higher frequency neighbors (filled triangles) is also shown. The inset bar graphs show simulated word recognition reaction times based on number of processing cycles required to reach the response threshold.




Figure 4. Results of Simulation 4: Word production. The left panel shows the time course of activation for the first phoneme in a high neighborhood density word (crosses) and a low neighborhood density word (squares). The inset bar graph shows simulated word production reaction times based on number of processing cycles required for the phoneme to reach the response threshold. The right panel shows the time course of activation for target words in high density neighborhoods (crosses) and low density neighborhoods (squares). Diamonds indicate activation of the neighbors during processing of the high neighborhood density target word.

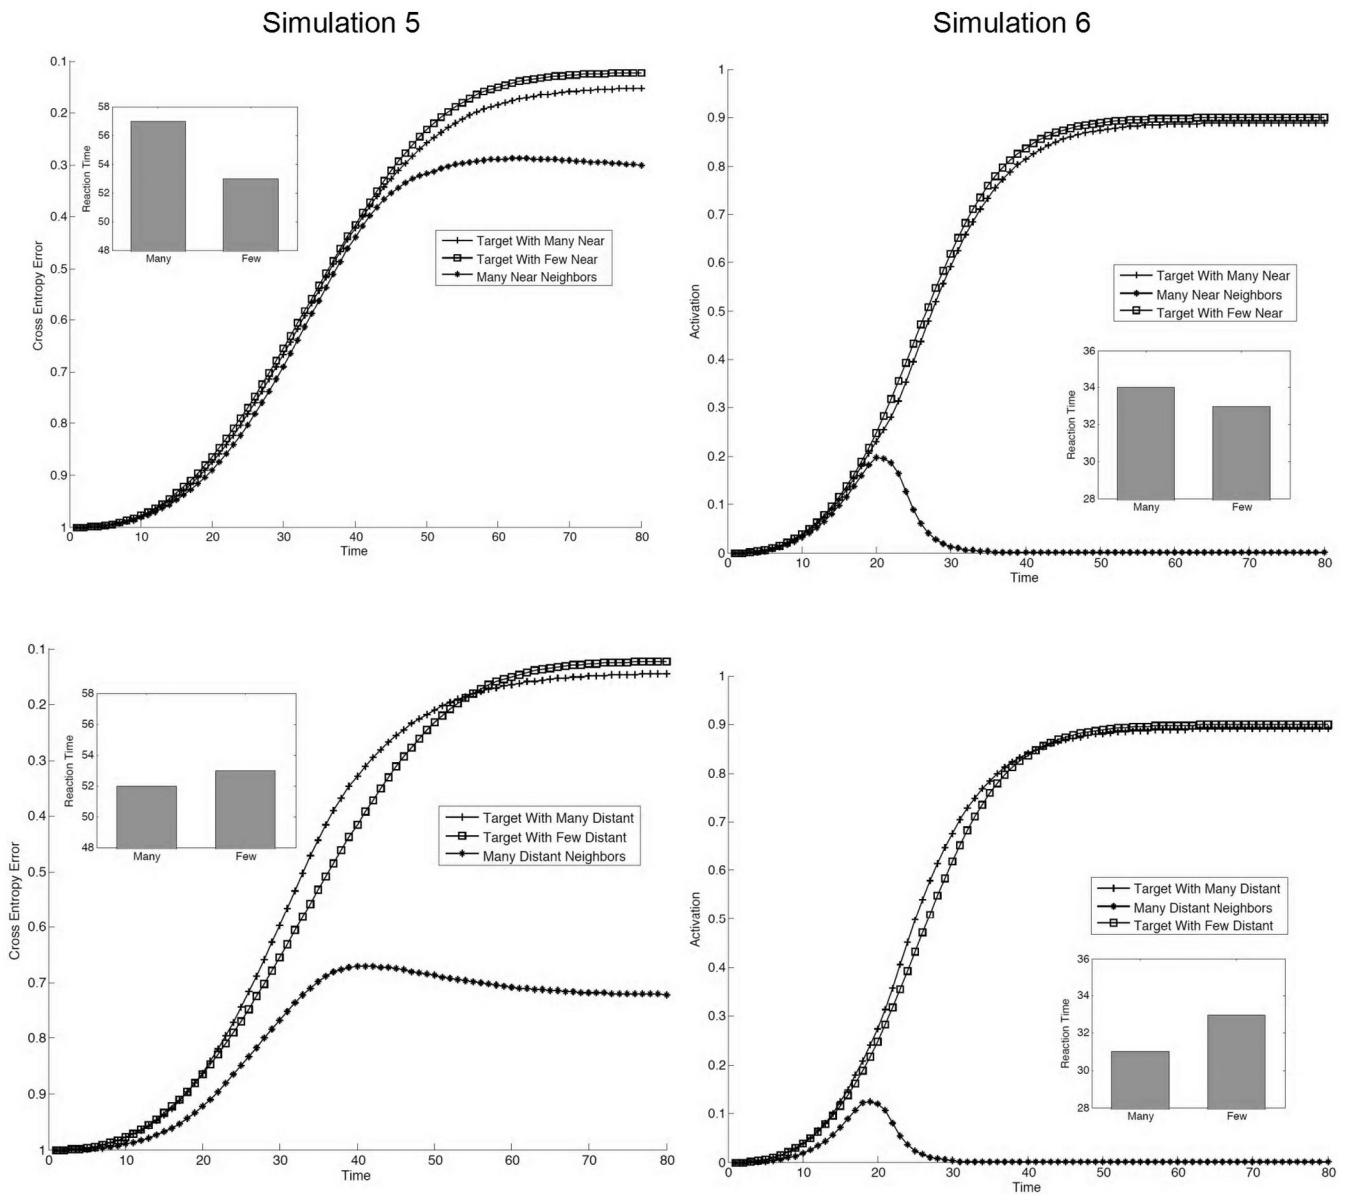


Figure 5. Results of Simulations 5 and 6: Effects of semantic neighbors on word recognition (Simulation 5, left panels) and word production (Simulation 6, right panels). Top row shows effects of near neighbors, bottom row shows effects of distant neighbors. The curves show the time course of processing (settling in Simulation 5, word activation in Simulation 6) for the low neighborhood density word (squares), the high neighborhood density word (crosses), and the neighbor words (circles). In the left panels, the vertical axis is reversed to reflect that lower nCEE corresponds to higher proximity to the target concept in semantic space. The inset bar graphs show simulated word recognition reaction times based on number of processing cycles required for the model to reach the response threshold.

Table 1

Effect of different kinds of neighbors in different tasks.

Neighbor/Task type	Behavior	Model
Form Neighbors		
Visual word recognition	Facilitation	Facilitation (Figure 3, left)
Visual word recognition with high frequency neighbor	Inhibition	Inhibition (Figure 3, middle)
Spoken word recognition	Inhibition	Inhibition (Figure 3, right)
Spoken word production	Facilitation	Facilitation (Figure 4)
Semantic neighbors		
Near neighbors	Inhibition	Inhibition (Figure 5, top)
Distant neighbors	Facilitation	Facilitation (Figure 5, bottom)