



Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2008 January ; 34(1): 65–79.

Attractor Dynamics and Semantic Neighborhood Density: Processing Is Slowed by Near Neighbors and Speeded by Distant Neighbors

Daniel Mirman and James S. Magnuson

Department of Psychology, University of Connecticut

Abstract

The authors investigated semantic neighborhood density effects on visual word processing to examine the dynamics of activation and competition among semantic representations. Experiment 1 validated feature-based semantic representations as a basis for computing semantic neighborhood density and suggested that near and distant neighbors have opposite effects on word processing. Experiment 2 confirmed these results: Word processing was slower for dense near neighborhoods and faster for dense distant neighborhoods. Analysis of a computational model showed that attractor dynamics can produce this pattern of neighborhood effects. The authors argue for reconsideration of traditional models of neighborhood effects in terms of attractor dynamics, which allow both inhibitory and facilitative effects to emerge.

Keywords

semantics; neighborhood density; word recognition; attractor networks

Semantic processing is a crucial, yet poorly understood, aspect of language processing. For example, there is general agreement that as a word is processed, related words and/or words with similar meanings are partially accessed or activated. However, the term *similar meanings* refers to very different relationships in different theories of semantic representation. The similarity structure of semantic representations goes to the very heart of knowledge and meaning—how knowledge is organized determines which concepts are similar or related and which ones are not. There are at least five approaches to semantic representation that propose strikingly different bases for semantic representations and consequently for the similarity relations among concepts. *Category-based hierarchical approaches* (e.g., Anderson, 1991; Collins & Loftus, 1975) define similarity in terms of category membership and location in a hierarchy. *Embodied cognition approaches* (e.g., Barsalou, 1999; Glenberg & Robertson, 2000) define similarity in terms of perceptual or action-based overlap. *Association-based semantic network approaches* define similarity in terms of connections defined by subject report (e.g., empirically collected association norms; Nelson, McEvoy, & Schreiber, 2004) or expert opinion (e.g., *Roget's Thesaurus*, Jarmasz & Szpakowicz, 2003; Steyvers & Tenenbaum, 2005). *Textual co-occurrence-based vector representations* (e.g., Landauer & Dumais, 1997; Lund & Burgess, 1996; Rohde, Gonnerman, & Plaut, 2004) define similarity in terms of cross-correlations of word co-occurrence. *Semantic microfeature-based representations* (e.g., McRae, Cree, Seidenberg, & McNorgan, 2005; Rogers & McClelland, 2004; Vigliocco, Vinson, Lewis, & Garrett, 2004) define similarity in terms of empirically determined feature overlap.

Despite concerted efforts over the last few decades, there is surprisingly little empirical basis for claiming that any one of these approaches captures the nature of human semantic representations. In part, this is because most studies of the structure of semantic knowledge have examined pairwise semantic similarity relations using paradigms such as priming, similarity ratings, and cued recall. This approach provides a measure of the similarity between any two concepts but risks missing the forest for the trees. An alternative, larger scope approach is to examine the effects of semantic neighborhood density (SND) or size—the number and/or proximity of neighboring representations—on word recognition. Analysis of neighborhood effects shifts the focus from the effect of a single related word on processing of the target word to the effect of the set of all similar words. Recent studies (Buchanan, Westbury, & Burgess, 2001; Siakaluk, Buchanan, & Westbury, 2003; Yates, Locker, & Simpson, 2003) taking this approach found that visually presented words in large or dense semantic neighborhoods were recognized faster than words in small or sparse neighborhoods.

Examinations of the effects of phonological and orthographic neighborhoods have a longer history. Written words with many orthographic neighbors (e.g., Sears, Hino, & Lupker, 1995) or phonological neighbors (Yates, 2005; Yates, Locker, & Simpson, 2004) are identified more quickly than words with few neighbors. In contrast, spoken words in dense phonological neighborhoods are identified more slowly than words in sparse neighborhoods (e.g., Luce, 1986; Luce & Pisoni, 1998). This contrast suggests that neighbors can have both facilitative and inhibitory effects, reflecting two opposing principles: Perceptual familiarity facilitates word recognition (i.e., items with many neighbors are more perceptually familiar, thus they are recognized more quickly), but competitor activation slows word recognition (i.e., items with many neighbors activate more competitors, and competition slows down recognition). Because spoken words are ambiguous at onset (e.g., /kæ/ could be the beginning of nearly 800 different words), competitor words with similar onsets can become active and compete for activation, thus (potentially) accentuating the competitive aspects of neighborhoods. In contrast, a visual word is presented all at once, so disambiguating information is available from the beginning, thus (potentially) reducing competition and accentuating the facilitative aspects of neighborhoods. Preliminary evidence from a study that forces subjects to read words letter-by-letter suggests that the difference between neighborhood effects in visual and spoken words results from parallel versus serial experience of words in the two domains (Magnuson, Mirman, & Strauss, 2007).

For orthographic and phonological neighborhoods, although individual metrics differ, there is an intuitive consensus that neighbors should be defined in terms of similarity of letters or phonemes (e.g., Luce & Pisoni, 1998; Sears et al., 1995). In contrast, the different approaches to semantics cited above give rise to radically different sets of primitives over which semantic distance can be defined. As a step toward developing a definition of semantic distance, Buchanan et al. (2001) compared two different measures of SND: one based on number of associates (derived from human participant generation of single associates to each target word; Nelson et al., 2004) and one based on mean distance to the 10 nearest neighbors according to the Hyperspace Analogue to Language (HAL) model (derived from co-occurrence statistics extracted from a large corpus of text, such that words that occur in similar contexts have similar representations and are close in semantic space; Lund & Burgess, 1996). Buchanan et al. found that the HAL-based measure was a better predictor of word recognition than the association-based measure and that the effect of SND was weaker for high frequency words. These results suggest that co-occurrence statistics capture neighborhood effects more accurately than semantic associates do (Balota, Cortese, Sergent-Marshall, Spieler, & Yap, 2004, also found weak effects of number of semantic associates). However, although association-based and co-occurrence-based approaches to semantic representation both tell us what concepts are similar (i.e., neighbors) and provide a measure of similarity, neither reveals why particular concepts are similar. For example, the associates of *car* include exemplars (e.g., *Toyota*), superordinate

terms (e.g., *transportation*), other vehicles (*truck, train, etc.*), interaction-related words (*drive, ride*), and various descriptive words (*fast, expensive, etc.*). It is intuitively clear that these are all related to *car*, so association norms are capturing relatedness (in fact, association norms and co-occurrence statistics provide excellent fits to behavioral word-pair similarity ratings; e.g., Rohde et al., 2004), but the various associates have radically different relationships to the target word, and therefore, the underlying similarity structure remains opaque.

In contrast to the opacity of association norms and co-occurrence statistics, representations based on semantic features explicitly encode the microstructure of semantic representations. Although feature norms do not capture the full complexity of semantic knowledge, they do capture a portion of that space at a level of detail that has a more transparent relationship to underlying similarity structure. Feature-based semantic representations are developed by asking human participants to generate features of a target concept (e.g., McRae et al., 2005, had subjects generate up to 10 features for many concepts; see also McRae, de Sa, & Seidenberg, 1997; Vigliocco et al., 2004). Models that represent semantic knowledge in terms of features provide powerful accounts of semantic priming (e.g., Cree, McRae, & McNorgan, 1999; Vigliocco et al., 2004), category-specific impairments (e.g., Cree & McRae, 2003), deterioration of semantic knowledge in progressive semantic dementia (Rogers et al., 2004), and speech errors and picture–word interference (Vigliocco et al., 2004). In general, feature-based models of semantic knowledge provide a coherent framework for understanding a very large set of phenomena (Rogers & McClelland, 2004). Crucially, with feature-based representations, similarity is defined by feature overlap, thus making the reasons for similarity explicitly available for analysis.

With respect to exploring neighborhood effects, feature-based measures of SND are potentially limited because the semantic neighborhood is strongly constrained to the items for which feature norms have been collected.¹ Association-based semantic neighborhoods do not have this limitation because an associate need not have been normed to be part of the target's semantic neighborhood (e.g., *type* can be an associate of *print* if a participant produces it, regardless of whether associates were collected for *type*; in a feature-based system, *type* and *print* could be neighbors only if feature norms have been collected for both). Co-occurrence-based semantic neighborhoods are limited by the size of the corpus, but because the corpus is typically very large (e.g., over 1 billion words), this is a very weak constraint. Thus, feature-based measures allow a finer-grain analysis of a smaller set of words relative to association-based and co-occurrence-based measures.

The present work addresses three questions: (a). How well do feature-based, association-based, and co-occurrence-based measures of SND capture semantic neighborhood effects? (b). Are SND effects facilitative (as previously demonstrated), inhibitory, or both? (c). Does a simple attractor model of semantic access capture patterns of neighborhood inhibition and facilitation consistent with the behavioral data?

In Experiment 1 we tested a large set of words in two word-recognition tasks and evaluated several measures of SND derived from feature-based, association-based, and co-occurrence-based semantic representations. The results suggested that, despite their limitations, SND measures based on feature representations are as good as those based on association norms and co-occurrence statistics at capturing the effects of semantic neighborhoods. Further, the results of Experiment 1 suggested a more complex story than simple facilitative or inhibitory effects. In Experiment 2 we specifically examined the independent effects of near and distant neighbors on semantic access and found that distant neighbors tend to have facilitative effects on semantic

¹Collecting and coding feature norms is a time-consuming task; the 541 item McRae et al. (2005) corpus is the result of a massive multiyear project and is the largest publicly available feature norm corpus.

access and near neighbors tend to have inhibitory effects on semantic access. Finally, analyses of settling rates in a simple attractor model of semantic access (Cree, McNorgan, & McRae, 2006) revealed that distant neighbors have early and transient facilitative effects and near neighbors have lasting inhibitory effects on settling, consistent with the human data. These results suggest that neighborhood effects can be understood by considering the specific impact of neighbors on attractor dynamics.

Experiment 1

The central goal of Experiment 1 was to examine the effects of SND in a relatively global manner. Specifically, we compared measures of SND derived from feature-based semantic representations to previously used measures derived from co-occurrence- and association-based representations. To provide the best basis for examining SND effects, we tested the 532 unique words from McRae et al. (2005) in lexical decision and semantic categorization (living thing judgment) tasks. In the following sections we first describe the six measures of SND that we tested and provide simple comparisons of the measures, then we describe the experimental methods and results, and finally we discuss how the results match and conflict with previous studies of SND and the implications of these results.

Measures of SND

Measures based on feature norms—The McRae et al. (2005) feature norm corpus contains 541 concepts (532 unique names) covering a broad range of living and nonliving concepts used in studies of semantic memory. Thirty participants from McGill University and/or the University of Western Ontario produced features for each concept. There were 2,526 unique features listed. Thus, each concept can be represented by a 2,526-element binary vector in which the elements code whether or not each feature was produced for each concept. These vectors are very sparse, ranging from 6 to 26 features for each concept ($M = 13.4$, $SD = 3.52$).

Measures 1–3 comprise our basic set of feature-based measures of SND. The first two measures are feature-based versions of a discrete number of neighbors measure (like number of associates) and a graded distance measure (like distance to neighbors in HAL); the third measure has been previously found to index the tightness of concept clusters (e.g., Cree & McRae, 2003). We also tested several embellishments of these basic measures (such as word frequency weighting, feature production frequency weighting, and feature distinctiveness weighting), but these embellishments did not improve the amount of variance captured by the measures.

1. *Number of Near Neighbors (NumNear)* is the number of concepts that have more than half of the target's features. The 50% threshold was chosen to balance between lower thresholds' tendency to eliminate the distinction between similar concepts and near neighbors and higher thresholds' tendency to eliminate all neighbors.² Note that this measure is asymmetric: If Concept A is a near neighbor to Concept B, Concept B is not necessarily a near neighbor of Concept A (i.e., if A has many features and B has few, and A and B overlap on half of B's features, the proportion of overlap will be high for B but lower for A). It is possible that semantic neighborhood relations are in fact asymmetric, which would be captured by this metric; symmetric similarity would be captured by another of our measures of SND.

²Because concepts typically have few features, a low threshold will count merely similar items as near neighbors. For example, *cheetah* has 13 features; at a threshold of 30% shared features, it has 28 near neighbors—a set that includes most of the mammals in the corpus, but at a threshold of 50% shared features, the near neighbor set is reduced to just the large predators. In general, as the threshold increases, the distribution of semantic neighborhood densities collapses towards 0 (i.e., at a sufficiently high threshold all concepts have no near neighbors).

2. *Mean Cosine (MeanCos)* is the mean cosine between target's feature vector and the feature vectors for every other item in the corpus. Cosine is a nonlinear measure of similarity that varies (for binary vectors) from 0 (no shared features) to 1 (identical feature vectors) and provides a symmetric measure of distance that reflects the similarity between active features in the target and neighbor.
3. *Proportion of significantly correlated feature pairs (Prop-CorrPairs)* is the proportion of feature pairs in the object's representation that tend to co-occur in the corpus (i.e., they share at least 6.5% of their variance; see McRae et al., 2005, for details). Objects with a greater proportion of co-occurring features should have nearer neighbors because their features tend to come in groups (i.e., they co-occur), and thus their feature overlap will tend to be greater. These correlations were based on vectors containing production frequency values (i.e., number of participants [maximum of 30] that produced this feature for this concept) rather than on binary feature vectors, which were used to compute NumNear and MeanCos. Production frequency feature vector versions of Num-Near and MeanCos produced virtually identical results, so the simpler measures are reported here.

Measures based on association norms. The University of South Florida free association norms (Nelson et al., 2004) contain associates produced for each of 5,019 target words. For each word, an average of 149 participants were asked to write the first word that came to mind that was meaningfully related or strongly associated to the presented word. As in previous studies (Balota et al., 2004; Buchanan et al., 2001; Yates et al., 2003), we defined the set of associates generated for each target as its semantic neighborhood.

4. *Number of Associates (NumAssoc)* is the number of associates for each target in the University of South Florida association norms. Previous studies found that words with more semantic associates were recognized more quickly (Balota et al., 2004; Buchanan et al., 2001; Yates et al., 2003). Because 125 of the McRae et al. (2005) items were missing from the University of South Florida association corpus, we excluded them from these analyses. Excluding these items from all analyses did not change the pattern of results, so we included all the McRae et al. items for all other analyses.

Measures based on co-occurrences. Recently, Rohde et al. (2004) proposed a semantic representation called the Correlated Occurrence Analogue to Lexical Semantic (COALS) that is similar to HAL (Lund & Burgess, 1996) in that it is also based on co-occurrence statistics from a very large corpus (1.2 billion word tokens representing 2.1 million word types taken from Usenet). COALS differs from HAL primarily in implementing a normalization technique to reduce the impact of high frequency closed class and function words and in ignoring negative correlation values when computing representation vectors (on the intuition that a target word's meaning is not informed by words that occur in very different contexts). COALS provides a better fit to behavioral word-pair similarity ratings and multiple-choice vocabulary tests (see Rohde et al., 2004), so we used COALS rather than HAL. We calculated semantic distance based on 500-element vectors (Rohde et al. used singular value decomposition to reduce the dimensionality of the vectors) using the Rohde et al. method: square root of the correlation between semantic vectors.³

5. *Mean Distance Within Set (COALS)* is the mean semantic distance to nearest 10 neighbors (cf. Buchanan et al., 2001) within the McRae et al. (2005) item set. A

³Correlation and cosine vector distance measures are equivalent, though the square root makes this distance measure different from the one used to compute MeanCos.

measure based on mean distance to all items in the McRae et al. set produced the same results, so we report the statistic most similar to the one used by Buchanan et al., who found that words with closer near neighbors (i.e., denser semantic neighborhoods) were recognized more quickly. One item (*dunebuggy*) was missing from the COALS corpus and was excluded from these analyses.

6. *Mean Distance Among All Items (COALS_all)* is the mean distance to nearest 10 neighbors within the 100,000 most common words in the COALS corpus.

Initial comparisons of SND measures—Before testing the SND measures, we explored two related questions about the measures themselves. (a). Were the neighbors identified by the different measures the same or different? (b). Was there evidence that the limited size of the feature-norm corpus limits semantic neighborhoods? Table 1 shows partial correlations among measures of SND after controlling for word frequency (HAL [Lund & Burgess, 1996] frequency norms, which Balota et al., 2004, found to be the best word frequency predictor of word recognition in young adults), length (in letters, phonemes, and syllables), bigram frequency, Coltheart's *N* (Coltheart, Davelaar, Jonasson, & Besner, 1977), and number of features. Computing partial correlations removes correlations among SND measures caused by control variables (e.g., two SND measures might be correlated because they are both correlated with word frequency) and allows the similarity between semantic neighborhood measures qua semantic neighborhoods to emerge. Not surprisingly, SND measures tend to cluster with measures based on the same underlying representations. One interesting deviation from this pattern is the strong negative correlation between MeanCos and PropCorrPairs, though both are positively correlated with NumNear. We examined more closely the neighborhoods formed by the different measures, separately for distance-based measures (MeanCos and COALS) and number-based measures (NumNear and NumAssoc).

Semantic neighborhoods defined by COALS_all were very different from semantic neighborhoods defined by MeanCos: These measures shared on average less than 1 of the 10 nearest neighbors. Semantic distance increased more rapidly across the nearest neighbors for MeanCos than for COALS_all. That is, according to the COALS_all measure, the second nearest neighbor was about 8.5% less similar to the target than the nearest neighbor; this decrease was 16% for MeanCos. For the 10th nearest neighbor (the farthest item included in the COALS_all neighborhood), this decrease was still only 25% for COALS_all but was 44% for MeanCos. For example, for *sheep*, the 10 nearest neighbors according to MeanCos are (with semantic distance in parentheses) as follows: *lamb* (0.60), *cow* (0.47), *goat* (0.44), *skunk* (0.32), *squirrel* (0.31), *fawn* (0.30), *donkey* (0.29), *otter* (0.29), *moose* (0.28), and *pig* (0.28); according to the COALS_all the 10 nearest neighbors are as follows: *bleating* (0.78), *goats* (0.75), *cows* (0.71), *cattle* (0.71), *herds* (0.65), *herders* (0.64), *oxen* (0.61), *goat* (0.60), *herd* (0.59), and *ruminants* (0.59). As this example demonstrates, the faster drop-off in similarity for feature-based measures could be due simply to neighbors that are missing from the feature norm corpus (e.g., about half of the COALS_all neighbors for *sheep* are not in the feature norm corpus); nonetheless, it demonstrates that feature-based representations are much more sensitive to differences in semantic similarity than co-occurrence-based measures. The very high positive correlation between COALS and COALS_all ($r = .6$) suggests that this is not merely an effect of the constraints of the McRae et al. (2005) corpus: When a COALS-based SND measure is computed from just the McRae et al. results, it produces nearly the same result as one computed based on the 100,000 most frequent words. Rather, it seems that feature-based representations are intrinsically more sensitive to semantic distance (we discuss possible reasons below).

The number of near neighbors (NumNear: $M = 0.9$, $SD = 2.3$) was consistently smaller than the number of associates (NumAssoc: $M = 13.5$, $SD = 5.0$). Only 8 of the items had more near neighbors than associates, and 273 items had 0 near neighbors (this large proportion of concepts

with 0 “near neighbors” severely limits the variability of this measure and puts it at a predictive disadvantage; however, as described below, this measure reveals an interesting and unique effect of near neighbors, thus validating the use of the relatively strict criterion). Of the 143 items that had both near neighbors and associates, for 43% (62 items) the near neighbors formed a subset of associates, for 27% (39 items) the near neighbors and associates were nonoverlapping sets, and for the remaining 30% (42 items) the neighbor–associate overlap ranged from 10% to 50%. In sum, it seems that for some items, feature-based semantic neighborhoods are very similar to associate-based semantic neighborhoods, but for other items the two semantic representations yield very different semantic neighborhoods. As with the distance-based measures, the smaller number of feature-based neighbors than associates could be due to neighbors that are missing from the feature-norm corpus and/or to feature-based measures’ greater sensitivity to semantic similarity.

In addition to corpus size, feature norms also differ in terms of the corpus contents. The McRae et al. (2005) feature norm database contains only basic-level concepts, such as “dog” and “chair”; thus the potential neighborhood is limited to other basic-level concepts, such as “cat” and “table,” respectively. In contrast, association norms and co-occurrence measures can, and do, produce neighborhoods that include parts (e.g., “legs” for *table*), features (e.g., “green” for *grass*), and categories (e.g., “pet” for *cat*) as neighbors. If semantic representations are structured such that basic-level terms have substantially different representations than features, parts, and so forth, then these different concepts would be distant neighbors, not near neighbors. This is another way in which feature-based representations may allow a distinction between near and distant neighbors that is obscured or missing from association-based and co-occurrence-based semantic representations.

These analyses comprise a preliminary comparison of the semantic neighborhoods defined by different representations of lexical semantics. They show that association-based, co-occurrence-based, and feature-based representations produce quite different semantic neighborhoods and that the limitations of feature-norm corpora may underestimate semantic neighborhood size and may reflect greater sensitivity to semantic similarity. Experiment 1 was designed in part to test whether the corpus size limitation undermines the ability of feature-based representations to capture SND effects.

Method

Participants—All participants were native English speakers and had normal or corrected-to-normal vision. The semantic categorization task was completed by 17 participants, and the lexical decision task was completed by 44 participants (22 participants randomly assigned to each half of the words). All participants were undergraduate students at the University of Connecticut who received course credit for participating.

Stimuli—The critical stimuli were the 532 unique orthographic forms in the McRae et al. (2005) feature norm database. These items were chosen to present the strongest test of feature-based representations. For the lexical decision test, 532 pronounceable nonwords were created that were matched in length to the words, and the items were divided into two lists of 532 items each (266 words, 266 nonwords) to keep the overall number of trials per participant equal for the two tasks and avoid fatigue effects.

Procedure—Stimuli were presented visually in 18-point black Courier font on a white computer screen background using E-Prime software (Psychology Software Tools, Inc., Pittsburgh, PA). We used a 17 in. (43.18 cm) cathode ray tube monitor with a refresh rate of 100 Hz and a resolution set to 800 × 600 pixels. Each trial began with a central fixation cross on screen for 1 s, then the stimulus was presented centered horizontally and vertically and

remained on screen until the participant responded (or until 5 s had elapsed). Each participant completed either the lexical decision version of the experiment or the semantic categorization version. These tasks were chosen to examine basic word-recognition accuracy and latency in one task that requires semantic access (semantic categorization) and one task that, in principle, does not (lexical decision). The experiment was completed in one session lasting approximately 20 min. Before beginning the critical block, participants completed a 20-trial practice session. In the lexical decision task, participants were asked to indicate whether each item was a word; in the semantic categorization task, the participants were asked to indicate whether each word referred to a living thing or a nonliving thing.

Results

None of the participants correctly identified *budgie* as a living thing, and only 13% correctly identified it as a word,⁴ so this item was excluded from analyses. In addition, trials on which reaction time (RT) was more than two standard deviations away from the overall mean were excluded from analyses (3.6% of lexical decision trials and 3.9% of semantic categorization trials). Overall accuracy was high in both tasks (lexical decision, 95.1% correct; semantic categorization, 94.0% correct), and mean RTs were 655 ms for lexical decision and 763 ms for semantic categorization. In the lexical decision task all critical items received “yes” responses, but in the semantic categorization task critical items included both “yes” (living) and “no” (non-living) responses; however, there was no difference in RT between “yes” (living) and “no” (nonliving) responses, $M_{\text{living}} = 761$ ms, $SD_{\text{living}} = 140$ ms; $M_{\text{nonliving}} = 764$ ms, $SD_{\text{nonliving}} = 112$ ms; $t(530) = 0.3$, $p = .76$; thus, the two types of items were combined in semantic categorization analyses (separate analyses are reported in Appendix A). Only trials on which a correct response was provided were included in the RT analyses.

The left section of Table 2 shows independent correlations between error rate and mean RT and control variables⁵ (top section) and measures of SND (bottom section). Not surprisingly, word frequency and length had strong correlations with both error rate and RT, particularly for the lexical decision task. Number of features (words with more features tend to be recognized more quickly, presumably because they have more robust semantic representations; see Pexman, Holyk, & Monfils, 2003) was also strongly correlated with error rate and RT in both tasks.

When the effects of word frequency, length, and number of features were partialled out, orthographic neighborhood (Coltheart’s N) no longer had significant correlations with error rate or RT for either task, and orthographic familiarity (bigram frequency) was significantly correlated only with lexical decision RT. Measures of SND also had significant correlations in one or both tasks (Table 3). This finding suggests that when the effects of word frequency and length are controlled, semantic neighborhoods play at least as large a role in visual word processing as orthographic neighborhoods do, and semantic neighborhoods are especially important when the task explicitly requires semantic access (semantic categorization relative to lexical decision). This conclusion must be tempered by the unique composition of our stimulus set, though weak or nonsignificant effects of orthographic neighborhood were also found in experiments testing a very large set of words (Balota et al., 2004).

⁴*Budgie* is a nickname for *budgerigar*, a small parrot that is a popular pet in Canada, where the feature norms were collected, but apparently unknown to University of Connecticut undergraduates.

⁵Two other control variables were examined: age of acquisition (AoA) and semantic congruency of orthographic neighbors (Pecher, Zeelenberg, & Wagenmakers, 2005; Rodd, 2004). Semantic congruency had no significant correlation with any dependent measures after word frequency was controlled, so it is omitted from analyses reported here. AoA captured unique variance, but it did not affect the variance captured by measures of SND. AoA measures (Gilhooly & Logie, 1980; Stadthagen-Gonzalez & Davis, 2006) are available for less than half of the McRae et al. (2005) words, thus including AoA as a control variable in the analyses would undermine the effort to provide a large-scale examination of SND effects. Because including AoA as a control variable did not affect the amount of variance captured by measures of SND, those analyses are not reported here.

The bottom right section of Table 2 shows partial correlations (controlled for word frequency, number of letters, number of phonemes, number of syllables, bigram frequency, orthographic neighborhood, and number of features) between error rate and RT and the measures of SND (analogous semipartial correlation results were virtually identical to partial correlation results). For the lexical decision task, NumAssoc and NumNear had marginal partial correlations with accuracy, and NumNear and COALS had reliable correlations with RT. The COALS correlation was negative, indicating faster word recognition in denser semantic neighborhoods, which is consistent with previous findings (Balota et al., 2004; Buchanan et al., 2001; Siakaluk et al., 2003; Yates et al., 2003). In contrast, the NumNear correlation was positive, indicating slower RT in denser semantic neighborhoods, a result that has not been found previously. For the semantic categorization task, MeanCos and COALS_all were significantly correlated with RT. COALS-based measures of SND, NumAssoc, NumNear, and PropCorrPairs were significantly correlated with semantic categorization error rate.

Discussion

Experiment 1 examined whether feature-based measures of SND can capture the effects of semantic neighborhoods and whether these effects are facilitative (as previously demonstrated) or inhibitory. Feature-based measures of SND captured as much RT variance in both lexical decision and semantic categorization tasks as association-based and co-occurrence-based SND measures did. Note also that in the case of ambiguous words (e.g., homophones) feature norms were collected on disambiguated words (e.g., *bat—baseball*), but association norms and co-occurrence measures do not make this distinction, just as the tasks used in this experiment did not disambiguate the stimulus words. Feature-based SND measures predicted unique variance in lexical decision and semantic categorization data despite this disadvantage and the limitations of the corpus. In general, the results indicate that feature-based measures of SND are at least as good as previously used association-based and co-occurrence-based measures of SND.

In general, SND effects emerged more clearly in RT than in accuracy measures, which is consistent with the high accuracy and consequently low error rate variability in both tasks. There was more variability in accuracy for semantic categorization than for lexical decision, though this variability was largely due to ambiguous items, such as foods. That is, many foods are ambiguous with respect to their status as a living thing (e.g., *corn, potato*), although even relatively unambiguously not-alive foods (e.g., *pickle, raisin, bread*) and some nonliving but highly living-related concepts (e.g., *beehive, shell* are natural kinds that house living things) also had low semantic categorization accuracy scores. These low semantic categorization accuracy items did not have systematic SND biases (the low-accuracy item set SND was less than 0.5 standard deviations from the overall mean SND according to all SND measures), but these task effects call for caution in interpreting our semantic categorization results. However, the qualitative similarity between semantic categorization and lexical decision results (e.g., significant partial correlations were all in the same direction for each measure) suggests that the semantic categorization results are not due to possible “living thing” task artifacts.

Our results replicate and extend previous studies of semantic neighborhood effects. Two of our SND measures (COALS and MeanCos) replicate previous findings of facilitative SND effects (Balota et al., 2004; Buchanan et al., 2001; Siakaluk et al., 2003; Yates et al., 2003). The COALS semantic representations are based on the same principles as HAL (which has previously been used to show facilitative SND effects), so this replication indicates that peculiarities of the items, tasks, or participants we tested cannot account for differences between our findings and previous studies. Specifically, the inhibitory effects of SND as measured by NumNear and COALS_all must be due to what those measures index, that is, differences with regard to which aspects of semantic neighborhoods are captured by the

different measures. The inhibitory/facilitative difference also cannot be due strictly to underlying semantic representations because the feature-based MeanCos measure showed facilitative effects, unlike NumNear (COALS and COALS_all show a similar reversal). One possible explanation for this reversal is that the NumNear measure captures a different aspect of semantic neighborhood structure than is captured by MeanCos.

One way that measures of neighborhood density may differ is in the impact of near neighbors versus distant neighbors. Semantic similarity distributions generally exhibit power law distributions, that is, concepts tend to have a few very near neighbors and many distant neighbors. One way to examine whether near or distant neighbors are the primary contributors to a particular measure of SND is to examine correlations between the measure of SND and number of neighbors according to the most liberal definition of semantic neighbor possible: items sharing at least one semantic feature. Because of the overwhelming number of distant neighbors, this liberal “number of neighbors” (different from NumNear, which is number of near neighbors) measure primarily reflects the number of distant neighbors. There was a strong positive correlation between number of neighbors and MeanCos ($r = .88, p < .001$), suggesting that the latter primarily reflects the number of distant neighbors. Because MeanCos is computed over the entire corpus, a concept with many additional (say, 100 more) distant neighbors (which will have low but nonzero cosine distances) will have a higher neighborhood density than a concept with a few additional near neighbors (which would have cosine distances of say, .75).

In contrast, a threshold-based measure such as number of near neighbors (NumNear) is not influenced by distant neighbors (the correlation with number of concepts sharing at least one feature was not significant, $r = .003, p = .95$); by definition, NumNear is most sensitive to the near neighbor structure. In sum, these analyses suggest that MeanCos and NumNear capture different aspects of semantic neighborhoods; specifically, NumNear reflects near neighbors, and MeanCos reflects distant neighbors. We found facilitative effects of MeanCos and inhibitory effects of NumNear, suggesting that near neighbors inhibit processing and distant neighbors facilitate processing.

For the COALS-based measures the near–distant distinction may also explain the paradoxical pattern of results. Co-occurrence vectors are defined over a very large corpus, and the McRae et al. (2005) corpus reflects only a very small sample of that space. Thus, it is possible that because the COALS measure was restricted to just the 10 nearest neighbors within the McRae et al. corpus, it captured the effect of relatively distant neighbors (in co-occurrence vector space), but the COALS_all measure captured the distance to the 10 nearest neighbors in the entire co-occurrence vector space and was thus more sensitive to near neighbors.

One way to test the hypothesis that near and distant neighbors have opposite effects is to examine correlations between semantic categorization and lexical decision RT data and number of near and distant neighbors. This type of analysis requires a continuous measure of similarity, so we used cosine distance between feature vectors. Defining “near” and “distant” neighbors requires a threshold, such that near neighbors would be those with cosine greater than the threshold, and distant neighbors would be those with cosine lower than the threshold but greater than zero (i.e., sharing some, but not many, features). We tested the correlations between RT in the two tasks, and near and distant neighborhood size was defined by three thresholds (.25, .5, and .75). Figure 1 shows that the results were consistent with an inhibitory effect of near neighbors (positive correlations) and a facilitative effect of distant neighbors (negative correlations, statistically reliable only in the semantic categorization task). We designed Experiment 2 to test the hypothesis suggested by these correlational results using a matched experimental manipulation of near and distant neighborhood size and a semantic task that avoids the pitfalls of living thing judgments.

Experiment 2

This experiment was designed to test the hypothesis that near semantic neighbors have an inhibitory effect and distant semantic neighbors have a facilitative effect on visual word processing. To test this hypothesis, we selected four sets of items and independently manipulated the number of near and distant neighbors while matching the items on all other criteria.

Method

Procedure—Experiment 1 results suggested that SND effects would emerge more strongly in a semantic task than in a lexical decision task but that items such as foods are ambiguous with respect to their status as a living thing. To avoid this problem while still using a semantic task, we used a concreteness judgment task in Experiment 2 (participants were instructed to indicate whether or not it was possible to touch the thing named on the screen). This task is less susceptible to ambiguity than the living thing task and all of the critical words referred to concrete things (see Table 4 for mean concreteness values for each condition; these values are on a 100–700 scale taken from the MRC Psycholinguistic Database; Wilson, 1988). In addition, there is some evidence that the use of a broad category makes for a more difficult semantic decision, thus encouraging more semantic processing and allowing semantic effects to emerge more clearly (Hino, Pexman, & Lupker, 2006). Consistent with this claim, in Experiment 1 we found stronger SND effects in the living thing judgment task relative to the lexical decision task, though the convergent results in the two tasks suggest that semantic neighborhood effects are relatively robust across tasks. The experiment began with 30 practice trials (with feedback; 15 concrete, 15 abstract trials) to familiarize participants with this task. Stimulus presentation details were the same as Experiment 1.

Materials—Twenty-five critical items were selected for each of four conditions: 2 (Many or Few Near Neighbors) \times 2 (Many or Few Distant Neighbors). Near neighbors were defined as having cosine greater than 0.5, distant neighbors were defined as having cosine less than .25 and greater than 0.0. For near neighbors, “many” was defined as at least 2, and “few” was defined as 0 or 1; for distant neighbors, “many” was defined as more than 200, and “few” was defined as less than 150. These thresholds were chosen because they divided the corpus of words into four relatively equal groups with very different neighborhood properties, and each of the groups was big enough to allow selection of a subset for matching on various control variables. This manipulation produced items with different near and distant neighborhood sizes as measured by cosine distance but equivalent semantic neighborhoods as measured by number of associates and mean COALS semantic density (see Table 4). In addition, conditions were matched on word frequency, familiarity, AoA (AoA data [Gilhooly & Logie, 1980; Stadthagen-Gonzalez & Davis, 2006] were available for only 30% of the words), length (in terms of letters, phonemes, and syllables), orthographic neighborhood (Coltheart’s N), number of semantically congruent orthographic neighbors, and number of semantic features, with approximately equal numbers of creatures, fruits and vegetables, and nonliving things in each condition. Control variables (including number of associates and COALS SND) did not differ reliably between conditions either in terms of analysis of variance main effects or interactions or in terms of pairwise comparisons. Condition means for critical and control variables are in Table 4, and the 25 words in each critical condition are in Appendix B. To balance the critical “yes” trials, we chose 100 filler words (“no” trials) from the MRC Psycholinguistic Database (Wilson, 1988) based on low (less than 400, values range from 100 to 700) concreteness and imageability scores ($M_{\text{concreteness}} = 295$; $M_{\text{imageability}} = 346$) and matched to the critical words on length and frequency.

Participants—Twenty-two undergraduate students at the University of Connecticut completed the experiment for course credit. All participants were native speakers of English and had normal or corrected-to-normal vision.

Results

The word *level* (few near, few distant neighbors) was judged by all participants not to be concrete, so it was excluded from analyses.⁶ Accuracy was high, and there were no significant effects of number of near and distant neighbors (means and standard errors are in Table 5; all $ps > .25$). Only trials on which a correct response was provided were included in the RT analyses. In addition, trials on which RT was more than two standard deviations away from the overall mean were excluded from analyses (4.1% of trials).

Figure 2 shows the RT results for the four critical conditions. As predicted, words with many near neighbors were categorized more slowly than words with few near neighbors (24 ms), $F_1(1, 21) = 17.3, p < .001$, partial $\eta^2 = 0.45$; $F_2(1, 95) = 4.6, p < .05$, partial $\eta^2 = 0.05$, and words with many distant neighbors were categorized more quickly than words with few distant neighbors (11 ms), $F_1(1, 21) = 5.2, p < .05$, partial $\eta^2 = 0.20$; $F_2 < 1$. There was no significant interaction between number of near and number of distant neighbors ($F_1 < 1$; $F_2 < 1$). The condition effects are within-subjects but between-items, and thus the items analyses have less power. We found no apparent outlier items (item means are in Appendix B), suggesting that the subjects effects (F_1) were not driven by a few outlier items and that the weaker effects by items (F_2) were indeed due to differences in statistical power.

Discussion

As in Experiment 1, our results replicate previous findings of facilitative effects of SND and provide new data casting those findings as part of a more complex pattern of SND effects. Our results suggest that previous findings of facilitative SND effects were primarily driven by distant neighbors. The fine grain of feature-based semantic representations revealed a novel finding: Near neighbors exert inhibitory effects on semantic access.

Traditional views of neighborhood density effects cannot account for these results because these results do not simply reflect competition (inhibitory effects) or familiarity (facilitative effects). Rather, the results reveal a complex interplay of neighbor distance and number such that inhibitory and facilitative effects occur simultaneously. We interpret these results in terms of an attractor model of semantic processing. The results suggest that distant neighbors create a gradient or gravitational force for faster settling into attractor basins, and near neighbors create conflicting subbasins and increase the likelihood of hitting a saddle point, which slows the completion of the settling process. In the following section we analyze an attractor-based computational model to test whether inhibitory effects of near neighbors and facilitative effects of distant neighbors could indeed be due to attractor dynamics.

Analysis of a Computational Model

To test the hypothesis that the inhibitory effects of near neighbors and the facilitative effects of distant neighbors could be due to attractor dynamics, we examined the settling patterns in

⁶In the McRae et al. (2005) feature norm collection study, all target items were concrete objects, which would encourage participants to list features for the concrete meaning of *level*. In the context of our concreteness judgment task, the abstract meaning, which is perhaps more common for college students, was activated. Because all participants responded that *level* was not concrete, including it in the analyses would merely distort the accuracy data, it would have no effect on the RT data because only correct response RTs are included. We also note that the conditions are equally well-matched on control variables (frequency, length, etc.) if *level* is removed from the list. *Level* was included in model analyses because this ambiguity does not affect model performance, though excluding it had no qualitative effect on model behavior.

an attractor model of semantic access (settling patterns were from the model reported by O'Connor, McRae, & Cree, 2006; see Cree et al., 2006, for model architecture and training details). O'Connor et al. (2006) trained the model to activate the appropriate semantic features for each of the 541 basic level concepts in the McRae et al. (2005) norms and for 20 superordinate category names (see O'Connor et al., 2006, for details). The input layer represented arbitrary "orthographic" input patterns (3 random units turned on among the 30 input units) and was fully connected to a semantic layer. The semantic layer consisted of one unit for each semantic feature, and each of the semantic units was connected to all other semantic units (with no self-connections). To compute a word's meaning, its word form was presented at the input layer, and activation propagated to the semantic layer for 20 ticks. To put the network in a neutral starting state that is generally consistent with most units being turned off (on average only 0.5% of feature units should be active), before a word form was presented, semantic unit activations were set to a random value between 0.15 and 0.25. At each tick, error was computed using the cross-entropy measure.

We tested the effect of near and distant neighbors on settling by examining the correlations between number of near and distant neighbors and cross-entropy error (CEE) at each tick. Model settling corresponds to decreasing error, thus, at every tick a positive correlation with neighborhood indicates that the neighbors are having an inhibitory effect (i.e., having more neighbors is associated with higher error), and a negative correlation indicates a facilitative effect of neighborhood (i.e., having more neighbors is associated with lower error). We computed correlations between normalized CEE⁷ and number of distant and near neighbors across all 541 concepts (similar to the analysis of Experiment 1 data, see Figure 1) at each tick. As for the behavioral data in Experiment 1, we used three definitions of "near" and "distant" neighbors based on cosine similarity thresholds. Near neighbors were defined as having minimum cosine thresholds .25, .5, or .75; distant neighbors were defined as having nonzero cosine less than .25, .5, or .75.

The results (Figure 3) showed a strong positive correlation with number of near neighbors that tends to increase over processing, consistent with an inhibitory effect of near neighbors (more near neighbors was associated with higher CEE). The correlation with distant neighbors was generally around 0 and not reliable, except for a clear dip to the negative side around tick 10 (reliable only with the stricter .25 maximum cosine threshold, which was used in Experiment 2). This dip is consistent with a transient facilitative effect of distant neighbors (more distant neighbors was associated with lower CEE) at an intermediate point in processing.

As a second test of the model, we examined model RTs for the words tested in Experiment 2. For behavioral tests it was critical to match the words on a variety of control variables such as word frequency, length, and so forth. However, this basic model is not sensitive to those properties, so we also tested the full set of words divided into the four design cells according to the thresholds used in Experiment 2 to define near versus distant neighbors and many versus few neighbors. The four sets of items were matched on the two factors that would impact model performance: number of features and concept familiarity. This approach produced 39 words with many near and many distant neighbors, 60 words with many near and few distant neighbors, 60 words with few near and many distant neighbors, and 92 words with few near and few distant neighbors, thus increasing the critical item set from 100 words to 251 words. Because the model analyses are restricted to the weaker by-items analysis (in Experiment 2 by-subjects analyses were stronger than by-items analyses), this increase in number of items is an important increase in power.

⁷Normalization (dividing raw CEE by the maximum CEE for each item) removes the effects of model starting state. Conceptually, normalized CEE corresponds to proportion of the distance from the target state (activation of all and only the correct features) to the starting state of the model.

We assume that behavioral response probability is some monotonic function of cognitive states related to cross-entropy error, that is, that participants respond when they have come sufficiently close in semantic space to the representation of the presented word. This was operationalized by setting a CEE response threshold and computing the number of ticks required for the model to come within the threshold normalized CEE level (this approach is typical for measuring RTs in attractor models [e.g., Cree et al., 1999], and more dynamic measures, such as minimum CEE reached and number of ticks required to reach minimum CEE, produced exactly the same pattern of results). We set the threshold at 0.2, but the model behavior is quite orderly, and a relatively large range of thresholds would produce the same qualitative pattern. The results for the Experiment 2 items are shown in the left panel of Figure 4 and the results for the full set of words are shown in the right panel of Figure 4. The model was consistent with the behavioral data: The model responded faster to words with fewer near neighbors (triangles lower than squares), Experiment 2 set, $F(1, 96) = 2.3, p = .13$; full set, $F(1, 247) = 25.8, p < .001$, partial $\eta^2 = 0.094$; and to words with more distant neighbors (lines angle up from left to right), Experiment 2 set, $F(1, 96) = 1.9, p = .17$; full set, $F(1, 247) = 4.1, p < .05$, partial $\eta^2 = 0.016$. There was no interaction (both F s < 1) between number of near and distant neighbors in the model data.

In sum, the model results demonstrate that an attractor model exhibits the contrasting inhibitory effect of near neighbors and facilitative effect of distant neighbors. These results suggest that the effect of neighbors on processing depends on their specific influence on an attractor surface that must be traversed in order to reach the target attractor, providing further support to attractor-based models as a promising approach to understanding semantic processing.

General Discussion

In Experiment 1 we found that feature-based SND measures capture unique variance (i.e., beyond seven control factors such as word frequency and length) in lexical decision and semantic categorization accuracy and latency. Despite the comparatively limited corpus for which feature norms are available, our feature-based measures performed as well as measures based on association norms and co-occurrence statistics did. Although each of the three approaches to semantic representation has strengths and weaknesses, feature-based representations provide a set of primitives, namely semantic microfeatures, for further exploration of the structure of semantic knowledge and the most transparent (albeit limited) theoretical link to the basis of semantic similarity.

The results of Experiment 1 showed an intriguing partial conflict with previous studies of SND: Some measures captured a facilitative effect of SND, consistent with previous results, but other measures captured an inhibitory effect of SND. Post-hoc correlational analyses suggested inhibitory effects of near neighbors and facilitative effects of distant neighbors. We tested this hypothesis directly in Experiment 2, in which near and distant neighborhood sizes were manipulated independently. The results of Experiment 2 confirmed that near neighbors tend to slow word processing and distant neighbors tend to speed word processing. We proposed that this complex pattern arises due to the consequences of near and distant neighbors on attractor dynamics: Distant neighbors create a gradient that facilitates settling to the correct attractor, and near neighbors are competitor attractors that delay word recognition.

As a test of this proposal, we analyzed the settling patterns of a simple attractor model of semantics trained on the 541 concept feature norm corpus. The results showed a relatively early transient facilitative effect of distant neighbors and a persistent and increasing inhibitory effect of near neighbors. This is precisely the pattern that would be predicted from the general attractor dynamics perspective that we proposed. Independent manipulation of near and distant neighborhood size revealed that the model does exhibit the inhibitory effect of near neighbors

and the facilitative effect of distant neighbors. These results suggest that attractor dynamics are a promising approach to understanding how neighborhood structure impacts word recognition because they capture the opposite effects of near and distant neighbors.

The opposite effects of near and distant neighbors mirror the opposing forces of familiarity facilitation and competitor inhibition and should apply to phonological and orthographic processing in addition to semantic processing. In speech perception researchers have also found both types of effects (Luce & Large, 2001), but they have attributed them to different processes or levels (Vitevitch & Luce 1999; see also Lipinski & Gupta, 2005). Specifically, Vitevitch and Luce (1999; also Luce & Large, 2001) argued that the inhibitory lexical neighborhood effects are due to competition at the lexical level and the facilitative phonotactic effects are due to sublexical benefits of familiarity. Both lexical neighborhood density and phonotactic probability capture aspects of phonological neighborhood density, but the measures are slightly different (though highly correlated) and seem to capture opposite effects, much like our different SND measures. Our analyses of semantic neighborhood effects suggest that both facilitative and inhibitory effects may arise from the same level depending on how distant and near neighbors impact attractor structure. Based on these results, we hypothesize that phonotactic probability measures may be capturing a facilitative influence of distant phonological neighbors, and lexical neighborhood density may be capturing an inhibitory influence of near phonological neighbors.

Orthographic neighborhood effects are generally thought to be facilitative (e.g., Sears et al., 1995), but there is at least one example of inhibitory orthographic neighbor effects: inhibition due to transposed-letter neighbors (e.g., Andrews, 1996). Words such as *salt*, which have a transposed-letter neighbor (*slat*) are processed more slowly than carefully matched words (*sand*). The standard definition of an orthographic neighbor is a word that differs from the target word by a single letter (e.g., Sears et al., 1995). The contrasting effects of transposed-letter neighbors and replaced-letter neighbors are consistent with the near/distant neighbor distinction under the hypothesis that transposed-letter neighbors (e.g., salt—slat) are nearer in orthographic representational space than replaced-letter neighbors (e.g., salt—sale). Future research will test this hypothesis directly.

The feature norm corpus used in the present experiments contained only concrete basic-level concepts, but the feature-based approach to semantic representation is not limited to just these types of words. For example, Vigliocco et al. (2004) developed a single feature-based semantic representational space for objects and actions, including both action verbs (e.g., *to scream*) and action nouns (e.g., *the trade*). Their key assumption was that the semantic representations of even relatively abstract words are grounded in our interactions with the environment and in partly modality-specific representations (see also Barsalou, 1999, for a comprehensive theory of grounding semantic knowledge in perceptual processes). Abstraction is correlated with a variety of semantic factors that affect neighborhood structure (e.g., number of features [Pexman et al., 2003], hierarchical structure [e.g., Breedin, Saffran, & Schwartz, 1998], and cross-correlation of feature pairs). These factors suggest that abstract concept representations are more sparsely distributed in semantic space, suggesting that facilitative distant neighbor effects should dominate. However, it is possible that because abstract concepts have vaguer representations (fewer features, possibly broader attractors), there is more potential for near neighbor interactions. Behavioral and computational tests are required to adjudicate between these conflicting hypotheses.

In sum, our results support two related points. First, among the different theories of semantic processing, those that can be cast as nonlinear dynamical systems characterized by attractor dynamics will be best able to capture the effects of neighborhood density, particularly the opposite effects of near and distant neighbors. In attractor models such as the one we analyzed

(Cree et al., 2006; O'Connor et al., 2006), these effects emerge naturally in the course of the settling process. It may also be possible to build the opposite effects of near and distant neighbors into models that do not have intrinsic nonlinear dynamic attractor properties (e.g., semantic networks based on association norms or categorical hierarchies), though we see this approach as ad-hoc and less parsimonious. Our results suggest that inhibitory effects of near neighbors and facilitative effects of distant neighbors are a property of attractor dynamics and would emerge in attractor-based models regardless of the underlying representation. Thus, we do not rule out association norms (or other semantic representations) as possible characterizations of semantic knowledge; rather we argue that these representations must be recast in attractor-dynamic terms in order to account for the present data. Second, traditional formulations of neighborhood effects—be they phonological, orthographic, or semantic—must be reconsidered in terms of nonlinear dynamic attractors, which allow both inhibitory and facilitative neighborhood effects to emerge without assigning them to different levels of processing.

Acknowledgements

This research was supported by National Institutes of Health Grants DC005765 to James S. Magnuson, F32HD052364 to Daniel Mirman, and HD001994 and HD40353 to Haskins Labs.

We thank Ann Kulikowski for her help with data collection and Nicole Landi, Jay Rueckl, Whitney Tabor, Ken McRae, and Lori Buchanan for their comments on an early draft. We are also grateful to Ken McRae, Doug Rohde, Douglas Nelson, and their colleagues for making their data publicly available and to Chris O'Connor, Ken McRae, and George Cree for providing their model data to us for analysis.

References

- Anderson JR. The adaptive nature of human categorization. *Psychological Review* 1991;98(3):409–429.
- Andrews S. Lexical retrieval and selection processes: Effects of transposed-letter confusability. *Journal of Memory and Language* 1996;35(6):775–800.
- 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]
- Barsalou LW. Perceptual symbol systems. *Behavioral and Brain Sciences* 1999;22(4):577–660. [PubMed: 11301525]
- Breedin SD, Saffran EM, Schwartz MF. Semantic factors in verb retrieval: An effect of complexity. *Brain and Language* 1998;63(1):1–31. [PubMed: 9642018]
- Buchanan L, Westbury C, Burgess C. Characterizing semantic space: Neighborhood effects in word recognition. *Psychonomic Bulletin & Review* 2001;8(3):531–544. [PubMed: 11700905]
- Collins AM, Loftus EF. A spreading-activation theory of semantic processing. *Psychological Review* 1975;82(6):407–428.
- 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, McNorgan C, McRae K. Distinctive features hold a privileged status in the computation of word meaning: Implications for theories of semantic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2006;32(4):643–658.
- 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.
- Gilhooly KJ, Logie RH. Age-of-acquisition, imagery, concreteness, familiarity, and ambiguity measures for 1,944 words. *Behavior Research Methods & Instrumentation* 1980;12(4):395–427.
- Glenberg AM, Robertson DA. Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language* 2000;43(3):379–401.

- Hino Y, Pexman PM, Lupker SJ. Ambiguity and relatedness effects in semantic tasks: Are they due to semantic coding? *Journal of Memory & Language* 2006;55:247–273.
- Jarmasz, M.; Szpakowicz, S. Proceedings of the Conference on Recent Advances in Natural Language Processing (RANLP-2003). Borovets, Bulgaria: 2003. Roget's Thesaurus and semantic similarity; p. 212-219.
- Landauer TK, Dumais ST. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review* 1997;104(2):211–240.
- 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.
- Luce, PA. Neighborhoods of words in the mental lexicon: Research on Speech Perception (Tech. Rep. No. 6). Bloomington: Indiana University, Speech Research Laboratory, Department of Psychology; 1986.
- Luce PA, Large NR. Phonotactics, density, and entropy in spoken word recognition. *Language and Cognitive Processes* 2001;16(56):565–581.
- Luce PA, Pisoni DB. Recognizing spoken words: The neighborhood activation model. *Ear and Hearing* 1998;19:1–36. [PubMed: 9504270]
- Lund K, Burgess C. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers* 1996;28(2):203–208.
- Magnuson, JS.; Mirman, D.; Strauss, T. Why do neighbors speed visual word recognition but slow spoken word recognition?. Paper presented at the 13th annual Conference on Architectures and Mechanisms for Language Processing; Turku, Finland. 2007 Aug.
- McRae K, Cree GS, Seidenberg MS, McNorgan C. Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods* 2005;37:547–559. [PubMed: 16629288]
- McRae K, de Sa VR, Seidenberg MS. On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General* 1997;126(2):99–130. [PubMed: 9163932]
- Nelson, DL.; McEvoy, CL.; Schreiber, TA. The University of South Florida free association, rhyme, and word fragment norms; *Behavior Research Methods, Instruments, & Computers*. 2004. p. 402-407. Retrieved from <http://w3.usf.edu/FreeAssociation/>
- O'Connor, CM.; McRae, K.; Cree, GS. Conceptual hierarchies arise from the dynamics of learning and processing: Insights from a flat attractor network. Proceedings of the 28th annual conference of the Cognitive Science Society; Mahwah, NJ: Erlbaum; 2006. p. 2577
- Pecher D, Zeelenberg R, Wagenmakers EJ. Enemies and friends in the neighborhood: Orthographic similarity effects in semantic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2005;31(1):121–128.
- Pexman PM, Holyk GG, Monfils MH. Number-of-features effects and semantic processing. *Memory & Cognition* 2003;31(6):842–855.
- Rodd JM. When do leopards get their spots? Semantic activation of lexical neighbors in visual word recognition. *Psychonomic Bulletin & Review* 2004;11(3):434–439. [PubMed: 15376791]
- Rogers TT, Lambon Ralph MA, Garrard P, Bozeat S, McClelland JL, Hodges JR, Patterson K. Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review* 2004;111(1):205–235. [PubMed: 14756594]
- Rogers, TT.; McClelland, JL. *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press; 2004.
- Rohde, DL.; Gonnerman, LM.; Plaut, DC. An improved model of semantic similarity based on lexical co-occurrence. Manuscript submitted for publication. 2004. Retrieved from <http://dlt4.mit.edu/~dr/COALS/>
- 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.
- Stadthagen-Gonzalez H, Davis CJ. The Bristol norms for age of acquisition, imageability, and familiarity. *Behavior Research Methods* 2006;38(4):598–605. [PubMed: 17393830]

Steyvers M, Tenenbaum JB. The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science* 2005;29(1):41–78.

Vigliocco G, Vinson DP, Lewis W, Garrett MF. Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology* 2004;48(4):422–488. [PubMed: 15099798]

Vitevitch MS, Luce PA. Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory & Language* 1999;40(3):374–408.

Wilson MD. MRC Psycholinguistic Database: Machine Readable Dictionary, Version 2. *Behavioral Research Methods, Instruments, & Computers* 1988;20:6–11.

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, 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 A

Supplemental Analyses of Experiment 1 Data

To examine differences in semantic neighborhood density (SND) effects for living and nonliving things, which received “yes” and “no” responses, respectively, in the semantic categorization, we computed separate partial correlations for living (199 concepts) and nonliving (332 concepts) things (see Table A1). The pattern of results closely corresponds to the overall data (see Table 2), but SND effects are much stronger for living things than for nonliving things (note also that statistical power is reduced by choosing smaller subsets of items; thus, some correlations that were reliable in the overall analysis are not reliable here). The difference in SND effects between living and nonliving things is generally not due to differences in SND variability—the standard deviation for each SND measure is very similar ($\pm 30\%$) for living and nonliving things (the one exception is NumNear [number of near neighbors], for which the standard deviation is more than five times greater for nonliving things than for living things). Critically, the opposing pattern of facilitative SND effects based on some measures (MeanCos [mean cosine] and COALS [Correlated Occurrence Analogue to Lexical Semantic, mean distance within set]) and inhibitory SND effects based on other measures (NumNear and COALS_all [mean distance among all items]) emerges here as well, although these effects appear to be largely driven by the living things.

Appendix B

Critical Words and Mean Reaction Times for Experiment 2

Many near neighbors				Few near neighbors			
Many distant neighbors		Few distant neighbors		Many distant neighbors		Few distant neighbors	
buffalo	599	guppy	872	rooster	655	hyena	796
chicken	530	squid	632	calf	655	coyote	655
fox	566	otter	688	pony	707	chimp	691
goose	630	owl	636	raccoon	615	dolphin	566
gopher	716	beaver	697	cow	621	cat	566
minnow	781	hawk	602	turkey	596	crown	645
moose	668	swan	632	rabbit	595	tripod	791
parakeet	737	dove	702	worm	676	saddle	695
pelican	642	cod	712	spider	566	axe	704
couch	644	church	739	rocket	675	raft	737
boat	668	bottle	618	cannon	684	level	NA
dagger	712	mittens	710	cigar	655	gown	781
gun	629	shawl	829	barrel	676	guitar	625
jet	636	jacket	657	barn	609	helmet	636

Many near neighbors				Few near neighbors			
Many distant neighbors		Few distant neighbors		Many distant neighbors		Few distant neighbors	
peg	747	shovel	705	cage	678	rope	666
pistol	603	sweater	635	tractor	689	basket	614
revolver	801	blender	751	sled	718	brick	593
sword	647	shoes	652	bridge	678	candle	678
tongs	734	coat	614	box	585	skirt	633
fridge	804	van	734	desk	660	closet	706
bucket	644	broccoli	641	olive	663	carpet	619
pin	757	cantaloupe	690	oak	609	lemon	627
nectarine	736	spinach	669	pumpkin	644	carrot	574
peas	666	orange	621	mushroom	654	pepper	637
plum	646	yam	667	garlic	632	corn	613




Figure 1.

Experiment 1 correlations between reaction time (RT) and number of near (black symbols) and distant (white symbols) neighbors based on three thresholds. Circles indicate results for the semantic (living thing) categorization (SC) task, triangles indicate results for the lexical decision (LD) task. * $p < .05$.




Figure 2. Experiment 2 reaction time (RT) results. RTs were slower for words with many near semantic neighbors (squares relative to triangles) and faster for words with many distant semantic neighbors. Error bars are one standard error.

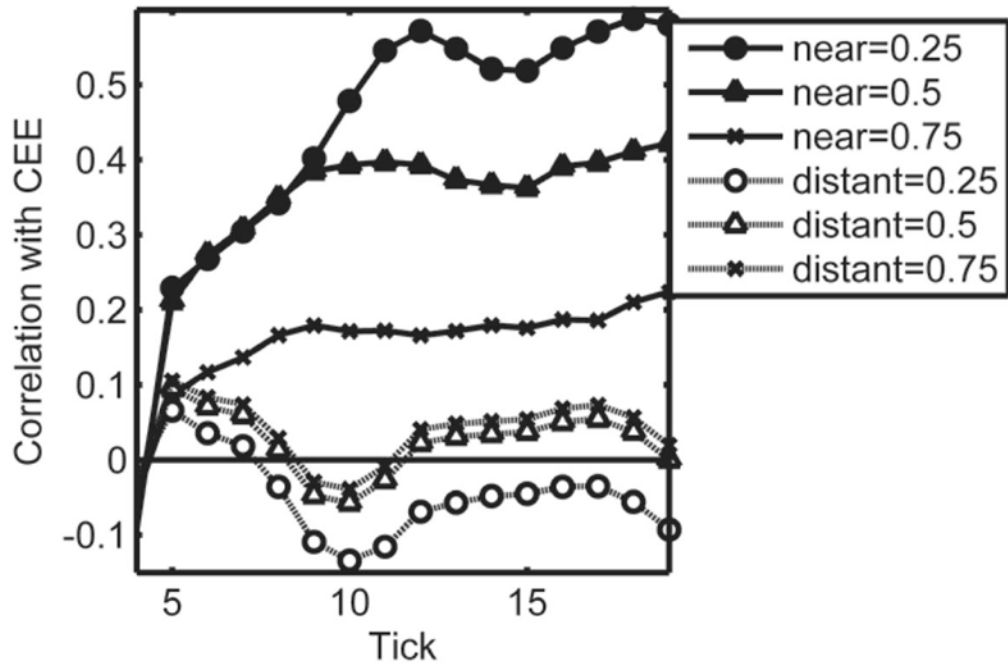


Figure 3. Correlations at each tick between cross-entropy error (CEE) and number of near neighbors (solid lines, black symbols) and number of distant neighbors (dotted lines, white symbols) across all concepts in the corpus. Correlations greater than $\pm .09$ are statistically reliable ($p < .05$). The correlations were computed for three levels of threshold defining near and distant neighbors. In the legend, “near = 0.25” indicates that near neighbors were defined as having minimum cosine 0.25 and “distant = 0.25” indicates that distant neighbors were defined as having maximum cosine 0.25 (for distant neighbors the minimum cosine was always 0).




Figure 4. Model response times (RTs; in ticks) to each of the four conditions tested in Experiment 2. The left panel shows RTs for the word sets used in Experiment 2. The right panel shows RTs for the full set of words based on the same threshold definitions of near and distant neighbors. Error bars are one standard error.

Table 1

Partial Correlations Among Measures of SND

SND measure	NumNear	PropCorrPairs	MeanCos	NumAssoc	COALS	COALS_all
NumNear	—					
PropCorrPairs	.39***	.39***	.30***	-.11**	-.06***	-.06***
MeanCos	.30***	—	-.23***	-.02	.19***	.17***
NumAssoc	-.11**	-.23***	.04	.04	-.11**	-.21***
COALS	-.06	.19***	-.11***	.12***	.12***	.03
COALS_all	-.06	.17***	-.21***	.03	.60***	.60***

Note. SND = semantic neighborhood density; NumNear = number of near neighbors; PropCorrPairs = proportion of significantly correlated feature pairs; MeanCos = mean cosine; NumAssoc = number of associates; COALS = Correlated Occurrence Analogue to Lexical Semantic, mean distance within set; COALS_all = mean distance among all items.

**
p < .05.

p < .01.

Table A1
 Error Rate and Mean RT Partial Correlations for the Six SND Measures in Lexical Decision (LD) and Living-Thing Semantic Categorization (SC) Separately for Living Things and Nonliving Things

Variable	Living things						Nonliving things					
	Error			RT			Error			RT		
	LD	SC		LD	SC		LD	SC		LD	SC	
Basic feature-based measures of SND												
NumNear	.11	-.06		.13*	.02		.17***	.02		.13**	.04	
MeanCos	.04	-.28***		.13*	-.21***		.05	-.07		.04	-.02	
PropCorrPairs	.08	.12*		.06	.14*		.02	-.02		.01	.00	
Association-based measure of SND												
NumAssoc	-.03	.11		-.17*	.11		.17***	.10		.05	.07	
Co-occurrence-based measures of SND												
COALS	-.09	.31***		-.27***	.06		.03	-.02		.02	-.04	
COALS_all	.03	.39***		-.09	.20***		.07	.10*		.00	.04	

Note. RT = reaction time; SND = semantic neighborhood density; NumNear = number of near neighbors; MeanCos = mean cosine; PropCorrPairs = proportion of significantly correlated feature pairs; NumAssoc = number of associates; COALS = Correlated Occurrence Analogue to Lexical Semantic, mean distance within set; COALS_all = mean distance among all items.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Table 2
Error Rate and Mean Reaction Time (RT) Correlations and Partial Correlations

Variable	Correlations				Partial Correlations			
	Error		RT		Error		RT	
	LD	SC	LD	SC	LD	SC	LD	SC
Control variable								
Ln(FreqHAL)	-.32***	-.20***	-.59***	-.39***	.08*	.08*	.09**	.03
No. letters	.03	.01	.33***	.21***	.01	-.03	.04	-.09**
No. phonemes	.05	-.02	.31***	.16***	.03	.10**	.03	.07
No. syllables	.13	.03	.34***	.19	.03	.10**	.03	.07
Bigram <i>f</i>	-.06	.02	-.24***	-.07	.03	.10**	.03	.07
Coltheart's <i>N</i>	-.10	-.02	-.30***	-.14***	.03	.10**	.03	.07
No. features	-.14	.02	-.23***	-.20***	.03	.10**	.03	.07
Basic feature-based measures of SND								
NumNear	.15***	.12***	.22***	.14***	.08*	.08*	.09**	.03
MeanCos	.07	.03	.11	-.02	.01	-.03	.04	-.09**
PropCorrPairs	.04	.12***	.06	.08*	.03	.10**	.03	.07
Association-based measure of SND								
NumAssoc	-.06	.02	-.20***	-.12**	.09*	.10**	-.04	.01
Co-occurrence-based measures of SND								
SND								
COALS	.03	.15***	.03	.05	.02	.11***	-.09**	.01
COALS_all	-.03	.12***	-.17***	-.01	.05	.19***	-.04	.12**

Note. LD = lexical decision; SC = living-thing semantic categorization; HAL = Hyperspace Analogue to Language; SND = semantic neighborhood density; NumNear = number of near neighbors; MeanCos = mean cosine; PropCorrPairs = proportion of significantly correlated feature pairs; NumAssoc = number of associates; COALS = Correlated Occurrence Analogue to Lexical Semantic distance within set; COALS_all = mean distance among all items.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Table 3

Partial Correlations for Measures of Orthographic and Semantic Neighborhood Controlling for Word Frequency, Length, and Number of Features

Measure	Error		RT	
	LD	SC	LD	SC
Orthographic measure				
Bigram <i>f</i>	.01	.03	-.11**	.03
Coltheart's <i>N</i>	-.02	.02	-.03	.06
Semantic measure				
NumNear	.08	.08*	.08**	.03
MeanCos	.01	-.03	.04	-.09**
PropCorrPairs	.03	.10**	.03	.06
NumAssoc	.08*	.10**	-.04	.01
COALS	-.01	.11**	-.09**	.00
COALS_all	.06	.19***	-.04	.10**

Note. RT = reaction time; LD = lexical decision; SC = semantic categorization; NumNear = number of near neighbors; MeanCos = mean cosine; PropCorrPairs = proportion of significantly correlated feature pairs; NumAssoc = number of associates; COALS = Correlated Occurrence Analogue to Lexical Semantic, mean distance within set; COAL-S_all = mean distance among all items.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Table 4
Experiment 2 Critical Item Condition Means (and SDs) for Critical and Control Variables

Variable	Many near neighbors		Few near neighbors	
	Many distant	Few distant	Many distant	Few distant
No. near neighbors	4.5 (3.0)	3.9 (3.1)	0.0 (0.0)	0.2 (0.4)
No. distant neighbors	227 (26.5)	109 (28.3)	249 (39.2)	113 (23.9)
No. associates	14.8 (5.2)	13.6 (2.8)	14.7 (6.3)	13.0 (5.2)
COALS_all	0.73 (0.072)	0.76 (0.11)	0.72 (0.08)	0.75 (0.078)
Ln(frequency)	1.8 (1.6)	1.9 (1.6)	2.3 (1.3)	2.5 (1.3)
Familiarity	5.0 (1.9)	5.6 (2.2)	5.8 (1.5)	5.9 (2.1)
Age of acquisition ^a	293.5 (63.9)	275.8 (61.2)	290.8 (77.2)	286.3 (31.0)
No. letters	5.4 (1.7)	5.4 (1.7)	5.2 (1.5)	5.2 (1.0)
No. phonemes	4.4 (1.6)	4.3 (1.6)	4.4 (1.4)	4.5 (1.1)
No. syllables	1.6 (0.8)	1.6 (0.6)	1.6 (0.5)	1.7 (0.6)
Coltheart's <i>N</i>	3.7 (4.3)	3.4 (4.4)	3.2 (4.1)	3.4 (5.0)
Congruent <i>N</i>	2.3 (2.8)	1.7 (2.8)	1.8 (2.5)	1.7 (3.1)
No. features	14.2 (3.9)	13.0 (3.4)	14.8 (3.2)	13.7 (3.9)
Concreteness	598.1 (26.0)	603.0 (17.2)	607.3 (19.0)	598.4 (16.9)
No. creatures	9	9	9	5
No. fruits/vegetables	3	5	5	4
No. nonliving	13	11	11	16

Note. There were 25 words in each condition. COALS_all = Correlated Occurrence Analogue to Lexical Semantic, mean distance among all items.

^aNote that age of acquisition data were available for only 30% of items (24%–36% for each condition).

Table 5
Experiment 2 Mean (and SE) Percent Correct Responses for the Four Critical Conditions

No. distant neighbors	No. near neighbors	
	Many	Few
Many	96.9 (1.0)	98.0 (0.6)
Few	97.3 (0.7)	97.6 (0.9)