

The Emergence of Semantic Meaning in the Ventral Temporal Pathway

Thomas A. Carlson^{1,2}, Ryan A. Simmons², Nikolaus Kriegeskorte³,
and L. Robert Slevc²

Abstract

■ In the ventral visual pathway, early visual areas encode light patterns on the retina in terms of image properties, for example, edges and color, whereas higher areas encode visual information in terms of objects and categories. At what point does semantic knowledge, as instantiated in human language, emerge? We examined this question by studying whether semantic similarity in language relates to the brain's organization of object representations in inferior temporal cortex (ITC), an area of the brain at the crux of several proposals describing how the brain might represent conceptual knowledge. Semantic relationships among words can be viewed as a geometrical structure with some pairs of words close in their meaning (e.g., man and boy) and other pairs more distant (e.g., man and tomato). ITC's representation of objects similarly can be viewed as a complex structure with some pairs of stimuli evoking similar patterns of activation (e.g., man

and boy) and other pairs evoking very different patterns (e.g., man and tomato). In this study, we examined whether the geometry of visual object representations in ITC bears a correspondence to the geometry of semantic relationships between word labels used to describe the objects. We compared ITC's representation to semantic structure, evaluated by explicit ratings of semantic similarity and by five computational measures of semantic similarity. We show that the representational geometry of ITC—but not of earlier visual areas (V1)—is reflected both in explicit behavioral ratings of semantic similarity and also in measures of semantic similarity derived from word usage patterns in natural language. Our findings show that patterns of brain activity in ITC not only reflect the organization of visual information into objects but also represent objects in a format compatible with conceptual thought and language. ■

INTRODUCTION

The domains of language and vision have been, for the most part, studied independently of one another, although there has been increasing interest in exploring the connections between these fields (Ferreira & Tanenhaus, 2007). An issue central to both fields is the organizational principles that underlie conceptual representations. In the domain of vision, this pursuit aims to understand how light patterns on the retina are transformed by the brain into meaningful units for cognition, for example, recognizing one's mother. In the domain of language, this pursuit aims to understand how communicative perceptual inputs are represented and decoded into meaningful content and how this content is encoded into linguistic form. An apparent bridge between these disparate domains is representation of conceptual meaning and semantic content. For us to recognize an object in the world and communicate this information to other language must interface with the contents of vision; To learn linguistic labels of visual forms, visual representations must be associated with semantic labels.

Semantic knowledge describes our ability to understand the meanings (i.e., semantic content) of these labels in different contexts, as well as the meanings of relationships between different labels. From these meanings and relationships, we derive the broader associations that constitute conceptual knowledge. In seeking to identify substrates for the connection between vision and semantic knowledge, inferior temporal cortex (ITC) is a well-suited candidate region. ITC has traditionally been considered to be the last exclusively visual area in the ventral visual pathway; (Logothetis & Sheinberg, 1996; Gross, 1992)—the so-called “what” pathway of vision (Ungerleider & Mishkin, 1982). In contrast to early visual areas that represent visual information as primitive features (e.g., color, orientation, and spatial frequency), which are unlikely to be the units of conceptual thought, ITC represents information at the “object” level of description, which often can be designated by basic level descriptions (e.g., face or chair). The organization of object representations in ITC by semantic category (Connolly et al., 2012; Konkle & Oliva, 2012; Reddy & Kanwisher, 2006; Downing, Jiang, Shuman, & Kanwisher, 2001; Kanwisher, McDermott, & Chun, 1997) also suggests that the way that we define objects, either explicitly or implicitly through our actions, relates to ITC organization (Mahon & Caramazza, 2011;

¹Macquarie University, Sydney, Australia, ²University of Maryland,

³Medical Research Council Cognition and Brain Sciences Unit, Cambridge, United Kingdom

Martin, 2007). Furthermore, recent neuroimaging studies have identified ITC as part of network of brain areas that encode semantic content (Huth, Nishimoto, Vu, & Gallant, 2012; Mitchell et al., 2008).

Although ITC was traditionally considered to be a visual area (Logothetis & Sheinberg, 1996; Gross, 1992), neuroimaging studies have shown that ITC activity is modulated by a wide range of tasks based on conceptual knowledge (for a review, see Martin, 2007), including nonvisual linguistic tasks such as auditory sentence comprehension (Rodd, Davis, & Johnsrude, 2005; Giraud, 2004; Davis & Johnsrude, 2003). Recent fMRI decoding studies have also shown coarse amodal category information (animal vs. tool) can be decoded from ITC activity (Simanova, Hagoort, Oostenveld, & Van Gerven, 2012). These findings suggest that the knowledge stored in ITC is not only visual but also conceptual. If so, patients with damage to ITC are predicted to show perceptual deficits in addition to accompanying conceptual/associative deficits. Indeed, whereas some patients with ITC damage exhibit deficits only on perceptual tasks (Basso, Capitani, & Laiacona, 1988; Silveri & Gainottib, 1988), many show equivalent impairments on both perceptual and conceptual/associative tasks (for a review, see Capitani, Laiacona, Mahon, & Caramazza, 2003). Caramazza and Shelton (1998), for example, presented a case with ITC damage with a specific deficit for animate objects, which was evident both in perceptual naming tasks, that is, “identify the animal shown in the image” and conceptual tasks, that is, “Is a cow a farm animal?” This dissociation is perhaps unsurprising if these tasks rely on knowledge represented in ITC because animacy is a predominant categorical boundary in the organization of ITC (Kriegeskorte et al., 2008; Kiani & Esteky, 2007). In fact, one might imagine that the role

of animacy in ITC might underlie the robust influence of animacy distinctions in language, including its role in grammar and in discourse (e.g., Dahl, 2008; Dahl & Fraurud, 1996).

In this study, we sought to examine whether the structure of semantic relationships among words is reflected in the representation of objects in ITC. If this is the case, then the organization of information in ITC will be reflected in our use of words. Our study made use of existing fMRI data from a study that used multivariate pattern analysis methods to measure the geometry of object representations in ITC (Kriegeskorte et al., 2008). We compared ITC’s geometry to the geometrical structure of semantic relationships derived from multiple measures: (1) explicit behavioral ratings of semantic similarity, (2) lexicographically motivated measures of semantic relatedness based on definitions (word senses) and hierarchical word relations (i.e., “is a”), and (3) emergent relatedness measures computed from distributional patterns of words in large text corpora. (Note that, for the purposes of this study, we use the terms “semantic similarity” and “semantic relatedness” synonymously.)

METHODS

Measurement of Visual Object Representations in Human Primary Visual Cortex and ITC

Our study made use of data from a previously published study (Kriegeskorte et al., 2008). Below we describe the relevant aspects of the study. For detailed methods, we refer the reader to the original article (Kriegeskorte et al., 2008).

The study by Kriegeskorte et al. (2008) characterized the geometry of object exemplar representations in the

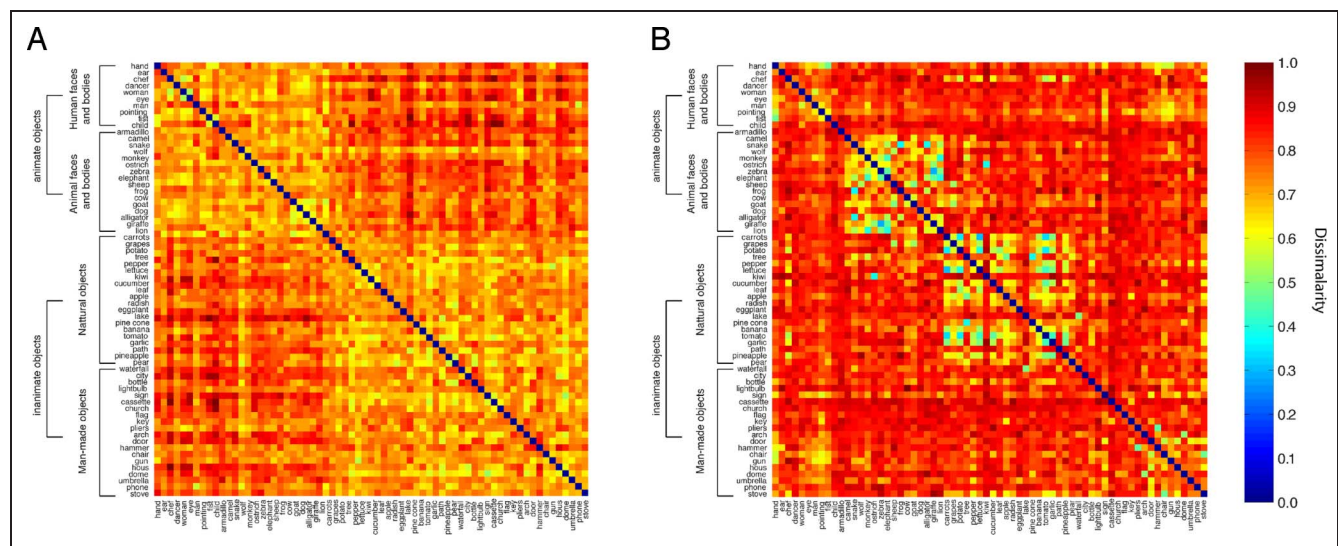


Figure 1. (A) Pairwise dissimilarity of stimulus exemplars’ neuronal activation patterns in ITC when viewing pictures of the stimuli. (B) Semantic dissimilarity of word labels associated with the stimuli evaluated using LSA. Cell color indicates dissimilarity value. High dissimilarity values indicate very distinct neuronal activation patterns/low contextual substitutability of words. Low dissimilarity values indicate similar neuronal activation patterns/high contextual substitutability of words.

human ITC and early visual cortex. Participants were shown 92 images of objects while their brain activity was recorded using fMRI. The interrelatedness of the representation of these stimuli within a brain area (e.g., ITC) can be construed as a geometrical structure in a high dimensional space, which quantitatively can be described as a dissimilarity matrix (DSM; see Figure 1A). Each entry of the DSM is a numeric value quantifying the “dissimilarity” between the brain activities for two object exemplars (e.g., an image of a man and a image of a tomato), where dissimilarity is computed as 1 minus the correlation between the two exemplars’ pattern of activation across voxels within the ROI. The complete DSM is all possible pairwise combinations of object exemplars.

Kriegeskorte et al. (2008) focused on two ROIs: ITC and early visual cortex. These ROIs were defined both anatomically and based on selectivity to the images used in the study. An ITC mask was defined manually in the functional slices, including all cortical voxels in the inferior occipito-temporal lobe including LO and extending anteriorly, but excluding early visual areas. Left and right ITC were defined similarly but constrained to the left and right hemisphere, respectively. An early visual cortex mask was manually defined in the slices as voxels in the calcarine sulcus in occipital cortex. To equate areas, our analyses were conducted on ROIs with equivalent numbers of voxels (316 voxels). The ROI used in the analysis was defined by selecting the voxels within each ROI that responded most strongly to object images, determined using an independent localizer. The number of voxels in each area was matched by adjusting the threshold.

Name Associations for Visual Object Stimuli

To study the relationship between the representational geometry of stimuli in the brain and the geometry of semantic relationships between labels associated with the stimuli, we first generated a set of object labels for the stimuli used in the study by Kriegeskorte et al. (2008). Twenty-five University of Maryland undergraduates participated in exchange for course credit. All participants reported speaking English as their native language. Before conducting the experiment, we removed nine of the images from the set of 92 images, as we a priori assumed these images would be given identical names in labeling. For example, the data set included multiple pictures of adult faces that would presumably all be labeled “face” or “human face.” Participants were shown the remaining 83 images in random order on a computer screen. For each image, the participants were asked to type the word that first comes to mind as a name for the picture. The most frequent name associated with each picture was chosen as that item’s semantic label. From this data, we identified additional images that were given the same label, for example, several of the pictures depicted different species of monkeys and were all labeled “monkey.” The images with overlapping labels were also excluded from the analysis. After removing the images

with identical labels, there were 67 exemplars with unique labels. Within this set, subject name agreement varied. There was 100% name agreement for 29 of the images, greater than 80% agreement for 47 of the images, greater than 60% agreement for 57 of the images, and greater than 50% for 63 of the images (all but four images). In the reported findings, we used all of the stimuli/labels regardless of the level of name agreement. We also examined the data using different cutoffs (e.g., using only images with greater than 60% name agreement), and the results were compatible with the findings using the entire stimulus set (data not shown).

Explicit Measures of Semantic Relatedness

The first measure of semantic similarity we used in our study was based on explicit ratings. We recruited a new group of twelve University of Maryland undergraduates to rate the semantic similarity of the labels. To complete the explicit semantic relatedness DSM, each participant in the study evaluated all 2211 word pairs (every pairwise comparison of the 67 words in our set) over three separate testing sessions. Participants were not asked to evaluate identities (i.e., the relatedness of a word with itself). Individual participants were shown a pair of words (labels) on each trial and asked to rate the semantic relatedness of the two words. Participants selected the degree of relatedness by positioning a bar on a GUI slider that returned a value between 0 and 100 (participants saw only the location of the bar on the slider and not the numerical value). Data were collected using custom software programmed in MATLAB (Natick, MD) using a Griffin PowerMate USB control knob.

The reliability of these explicit judgments was assessed using intraclass correlation (two-way, consistency, average-measures ICC; Hallgren, 2012), yielding an ICC coefficient of 0.946 (95% CI = ± 0.003). Relatedness was thus rated similarly across participants, suggesting a minimal amount of measurement error and reasonable statistical power for subsequent analyses.

Measures of Linguistic (Dis)similarity

Our study also examined several computational measures of semantic relatedness. For each measure, all pairwise combinations of word labels were compared and used to construct a DSM. Figure 1B shows the DSM for latent semantic analysis (LSA; Landauer & Dumais, 1997), one of the measures in our study. Our study used several measures from two general methods: one based on explicit hierarchical category structure (WordNet based measures) and one based on statistical patterns of occurrence in large corpora of text (LSA and Correlated Occurrence Analogue to Lexical Semantic [COALS]).

WordNet Measures

WordNet (Fellbaum, 1998; Miller, 1995) is a hierarchical lexical database that represents words by their dictionary

definitions (glosses), their part of speech (nouns, verbs, adjectives), and by membership in “synsets,” defined as sets of synonyms that are interchangeable in some context. Note that words of different parts of speech occupy their own hierarchical spaces without connecting nodes, and so most WordNet measures are unable to compare across different parts of speech. All items used in this study were nouns, so this limitation is not relevant to our evaluations. In WordNet, individual word senses are connected to each other in the hierarchy through hyper/hyponymic relationships (i.e., “X is a Y”), and synsets are connected to each other via a variety of relations (e.g., metonymy, holonymy, etc.). The top level of the hierarchy consists of abstract root nodes, which may be subdivided into additional nodes (e.g., the primary root node for nouns is “entity,” which is divided into “animate,” “inanimate,” “composition,” and “roles.”). Similarity between items in the WordNet database was calculated using three different methods implemented in the WordNet::Similarity Perl module (Pedersen, Patwardhan, & Michelizzi, 2004). For our calculations, the first sense of each word type was consistently chosen as the representative token. Although relatedness values vary by relative word sense, WordNet senses are ranked by frequency, so the first sense is more likely to be the “correct” token. Previous studies have indicated this to be a reasonably precise heuristic (Hawker & Honnibal, 2006; Moldovan & Novischi, 2004).

The WordNet PATH Measure

The most straightforward of the WordNet similarity rubrics is the PATH measure, defined as the multiplicative inverse of the shortest distance between two word senses in the hierarchy. The distance is calculated as the number of “steps” or “nodes” it takes to get from one sense to another in the hierarchy, using hyper/hyponymic relationships. Thus identical senses have a path distance of 1, and as distance between senses increases, the value of the path measure decreases toward 0.

The WordNet LESK Measure

The PATH measure uses the explicit synset links delineated by the WordNet hierarchy to calculate relatedness. Other WordNet measures use glosses of words encoded in WordNet under the assumption that related words or concepts will use similar words in their glosses. For example, “bar” and “drink” are not closely connected through “is-a” steps and relationships (PATH relatedness = 0.1111). However, the concepts share an obvious conceptual association, which is at least in part reflected in the similarity of their respective glosses: “a room or establishment where alcoholic drinks are served over a counter” and “a single serving of a beverage.” When comparing synsets using the extended gloss overlap measure (or LESK), WordNet

will search the glosses of the immediate “neighbors” of the target words (that is, other concepts connected to the target through a single hierarchical “step”). Overlaps scores are additive, and consecutive words or phrases are scored exponentially; if two glosses share the word “paper” the LESK score is 1, but if they share “prepared paper” the LESK score is 4, “specially prepared paper” is 9, etc. The relatedness between two concepts is calculated by summing the LESK scores for all glosses and normalizing by the size of the glosses (Banerjee & Pedersen, 2002, 2003).

The WordNet VECTOR Measure

The final WordNet measure we implemented was gloss vectors, which, like LESK, is calculated by the content of the glosses encoded in WordNet. A context vector is constructed for a particular word as the resultant of the co-occurrence vectors for each word in its gloss. The values of a co-occurrence vector are determined by the frequency with which a given word occurs with another throughout the WordNet corpus. The vector for “bar” is the centroid (or sum) of the vectors for room, establishment, alcoholic, drink, serve, and counter (conjunctions and articles are excluded), whereas the vector for “drink” is the centroid of single, serving, beverage. The relatedness between concepts is calculated as the cosine of the difference between their context vectors, measuring the relative divergence of the vectors in a hypothetical semantic space (Patwardhan & Pedersen, 2006).

Distributional Measures

Semantic relatedness can also be computed based on emergent patterns of co-occurrences in text. These distributional measures have generally been found to correlate well with explicit human judgments of semantic relatedness (Boyd-Graber & Fellbaum, 2006; Rohde, Gonnerman, & Plaut, 2005). However, it is important to note that these patterns do not directly measure word meaning and may not encode certain important aspects of meaning (Glenberg & Mehta, 2008). Rather, they use the patterns in which we use words in regular linguistic context as a model for implicit semantic structures and relationships (see Discussion). In our study, we examine two distributional measures of semantic relatedness: LSA and COALS.

LSA (Landauer & Dumais, 1997) is based on the assumption that words that mean similar things will tend to occur in similar contexts. Similarity is measured through second-order co-occurrence. First-order co-occurrence describes words that appear together in a particular context, whereas second-order co-occurrence describes the relationship between words with shared first-order co-occurrences. For example, “tire” and “windshield” may not appear together in a sentence, but both will appear in texts with words such as “car,” “drive,” “highway,” etc.

As with the VECTOR measure, each word is represented by a vector of co-occurrence values; instead of gloss definitions, LSA calculates co-occurrence over a large corpus of text. We used the TASA (Touchstone Applied Science Associates) college reading level corpus of 92,409 word tokens across 37,651 documents, calculated through an on-line interface (lsa.colorado.edu). Because of the high dimensionality of the corpus, vectors are factored using singular value decomposition (employing 419 factors). Relatedness between words is represented by the cosine of the difference between co-occurrence vectors. The relative divergence of the vectors is taken as a measure of the degree of contextual substitutability of the words in natural language.

COALS (Rohde et al., 2005) is another measure based on co-occurrence patterns. As with LSA, vectors of first-order co-occurrences are constructed for each target word; however, instead of calculating co-occurrences across an entire document, COALS uses a ramped “window” size of 4. Only the four most proximal words on either side of the target word are used to create the vector, with the closest receiving the highest scores. Scores were calculated using an on-line interface (dlt4.mit.edu/~dr/COALS), using a corpus of 1.2 billion word tokens across 9 million distinct documents gathered from Usenet postings. As with LSA, dimensionality reduction is employed via singular value decomposition (in this case, using 800 factors). COALS calculates the conditional rate of co-occurrence (that is, does word X occur more or less often in the vicinity of word Y than its average across documents) by computing Pearson’s correlation coefficients be-

tween constructed word vectors. These coefficients are normalized by setting all negative values equal to 0 and taking the square root of the positive values.

Statistical Comparisons

Comparisons between DSMs were conducted in the representational similarity analysis framework (Kriegeskorte et al., 2008). For each comparison (e.g., explicit ratings and human IT), we computed a nonparametric Spearman’s rank correlation between entries of the two DSMs. In the analysis of the complete DSMs (shown in Figure 2), the entries in the upper right triangle of two DSMs are correlated with one another. Note DSMs are symmetrical so the lower left triangle is identical to the upper right. In the fine-grained analysis (shown in Figure 3), only the relevant entries (e.g., comparisons “within” category) are correlated.

To evaluate statistical significance, we compared the actual correlation to a null distribution of correlation values. To generate the null distribution, we randomly shuffled the labels in the DSMs and computed the correlation between the (shuffled) DSMs. This was repeated 10,000 times to generate the null distribution. The reported *p* values are percentage rank of the actual correlation value within the null distribution.

RESULTS

The broad aim of our study was to test whether the semantic structure of linguistic labels is reflected in the

Figure 2. Comparisons between DSMs constructed using explicit human judgments of word label similarity (left), computational measures of semantic similarity (middle column), and neuronal activation patterns in primary visual cortex and ITC (right column). Lines connecting the DSMs indicate the comparisons. Numbers above each line are the Spearman’s rank order correlation for the comparison. Dotted lines indicate comparisons that are not statistically significant. Filled lines are those comparisons that are statistically significant ($p < .01$, bootstrap test) with the thickness of the line indicating the strength of the correlation. The four boxes superimposed on the human judgment DSM indicate the qualitative category clusters apparent in semantic similarity measures.

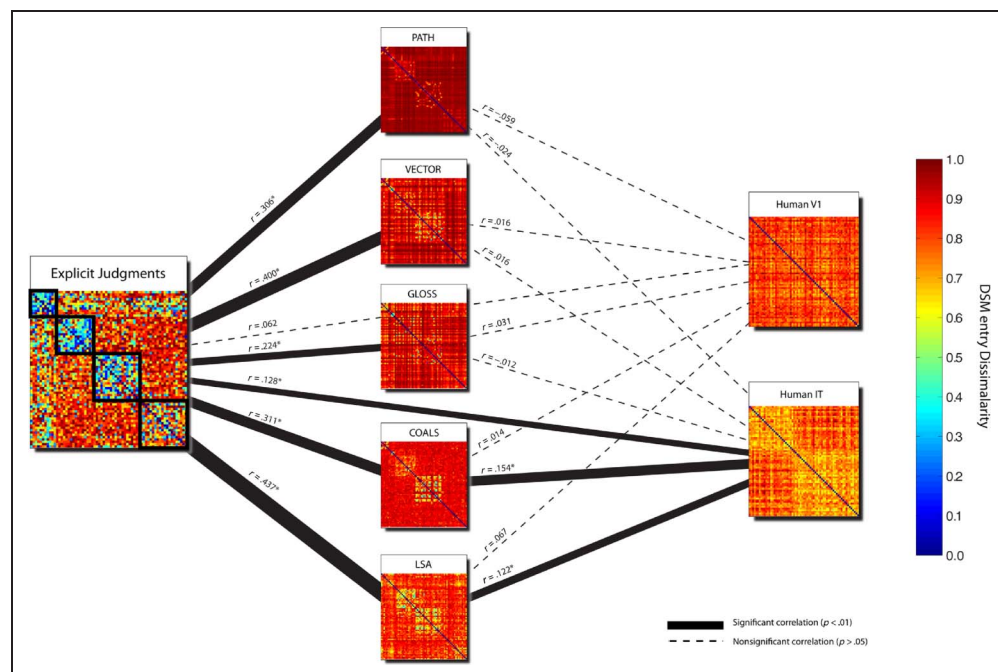
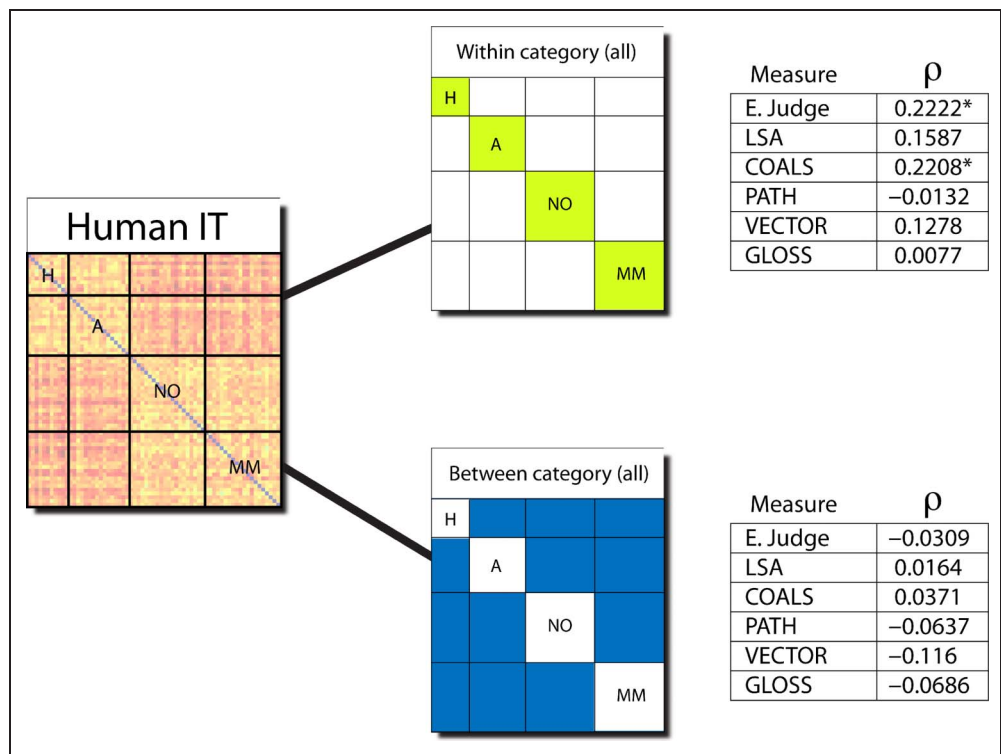


Figure 3. Analysis of within- and between-category substructure. Shown is the human ITC data divided diagrammatically into the four-category substructure observed in the analysis of semantic relatedness (H = humans, A = animals; NO = natural objects; MM = man-made objects). In the substructure analysis, the relevant DSM entries (within-category comparisons shown in green and between-category comparisons shown in blue) are analyzed separately. Next to the figures are tables showing the Spearman's rank order correlations for each of the measures. Asterisks denote significance $p < .01$ (bootstrap test).



brain's representation of objects and specifically the representation of objects in ITC. We did this using representational similarity analysis (Kriegeskorte et al., 2008) to compare representations (quantified as DSMs) derived from multiple lexographically based and distributionally based measures of semantic relatedness to brain activity in human primary visual cortex and human ITC (see Methods). Figure 2 summarizes our findings for primary visual cortex and human ITC. We additionally analyzed left and right ITC separately. The results for the two hemispheres did not depart substantially from the findings of bilateral ITC. Below, we only present the data for bilateral ITC. The data for left and right ITC are given in Supplementary Figure 1.

Explicit Behavioral Ratings of Word Similarity Correlate with Distributional and WordNet Measures of Semantic Similarity AND with Activity Patterns in ITC

WordNet measures of semantic relatedness are explicitly implemented in the database, whereas distributional measures determine relatedness based on statistical regularities of word usage in text corpora. To validate these two types of measures, we first examined their relationship to explicit ratings of semantic similarity. For each measure, we constructed a DSM and evaluated whether participants' explicit ratings (leftmost DSM in Figure 2) corresponded with the WordNet and distributional measures (DSMs shown in the middle column of Figure 2). In each of the comparisons, we found a significant correla-

tion ($p < .01$), thus showing that both the WordNet and distributional measures correspond well with explicit evaluations of semantic relatedness.

Our hypothesis predicts that the ratings of semantic similarity will also be reflected in the brain's representation in ITC. We further predicted there would be no relationship between primary visual cortex and semantic similarity, as early visual cortex represents visual stimuli in terms of primitive image features (color, contrast, edge orientations, etc). The DSMs for primary visual cortex and ITC are shown in the rightmost column of Figure 2. Concordant with our predictions, we found a significant correspondence between the participant ratings and the ITC representation ($p < .01$), but not the representation in visual cortex ($p > .05$).

Our initial analysis validates each of the measures of semantic relatedness used in the study, both those theoretically motivated and those emergent from word usage, by showing these measures correlate with explicit ratings. Notably, we also found support for our hypothesis by showing a correspondence between explicit ratings and the brain's representation of objects in ITC.

The Structure of Object Representations in IT Does NOT Correlate with WordNet Measures

WordNet describes the semantic relationship between words within a hierarchical framework. WordNet's structure emphasizes two relational principles: hierarchical super-subordinate relations and synonymy (similarity in meaning). Kriegeskorte et al. (2008) observed hierarchical

structure in ITC object representations. WordNet thus is a natural starting point to examine the relationship between the structure of semantic knowledge and the structure of object representations in ITC.

We examined whether each WordNet measure corresponded with the representation of objects in visual cortex and ITC (see Figure 2). Not surprisingly, none of these measures correlated with the brain's representation of objects in visual cortex ($p > .5$ for all comparisons). Somewhat surprisingly, we also found no correspondence between any of the WordNet measures on the brain's representation in ITC ($p > .05$ for all comparisons). The lack of relationship could be attributed to two related explanations. First, WordNet's hierarchy may be mismatched with ITC's hierarchical organization, especially with respect to the PATH measure, which explicitly encodes WordNet's semantic hierarchy. It is notable, however, that WordNet's glosses, as utilized by the VECTOR and LESK measures, also tend to encode a hierarchical structure, for example, the gloss for "cat" is "a feline mammal." The hierarchical mismatch explanation was supported by qualitative comparisons of the ITC and the WordNet DSMs. The WordNet DSMs organize into three or four clusters: human, animal, natural objects, and possibly a fourth category of man-made or artifactual objects (highlighted on the DSM for explicit judgments). In ITC, the DSM has two bright square regions corresponding to the locations of animate and inanimate exemplars (see Figure 1A). These two distinct regions indicate that a fundamental organizational principle of ITC is animacy (Kriegeskorte et al., 2008), although it should be noted that Kriegeskorte et al. (2008) also found that ITC exhibited a less apparent human/animal distinction. This hierarchical mismatch might explain the weak link between the WordNet measures and ITC. This explanation, however, is unsupported by the explicit ratings. The explicit ratings DSM also shows the three/four-way clustering observed in the WordNet DSMs, yet the structure of explicit ratings did match with ITC. Alternatively, it is possible that WordNet's glosses and forced hierarchical structure, which are based on an explicit hypothesis of what semantic knowledge is and how it is organized, might ineffectively capture "natural" semantic relations. If so, it may be more appropriate to consider measures of semantic relatedness that emerge from distributional patterns of word usage in natural language.

The Structure of Object Representations in ITC Matches with Distributional Measures of Semantic Structure in Language

We next examined whether the structure of object representations in ITC correspond to the emergent structure of word meaning arising from statistical occurrence patterns of words in large text corpora. The most well known measure of this type is LSA (Landauer & Dumais, 1997), which evaluates semantic relatedness based on the similarity of

text environments in which words tend to occur. In addition to LSA, we also examined COALS (Rohde et al., 2005), which evaluates semantic relatedness based on the (conditional) co-occurrence of words with each other (rather than the extent to which words occur in similar environments, as in LSA).

We compared the brain's representation in primary visual cortex and ITC to LSA and COALS (see Figure 2). For both measures, we found no evidence of a correspondence with visual cortex ($p > .05$ for both comparisons), as expected, and a match with ITC ($p < .01$ for both comparisons). These findings indicate that the geometry of the brain's representation in ITC matches with our patterns of word use in written language. In the context of the WordNet data (described above) these findings support the idea that distributional measures of semantic similarity like LSA and COALS may capture different aspects of semantic organization than WordNet measures (cf. Maki & Buchanan, 2008). Broadly, these results show that emergent structure from text corpora better reflect both behavioral judgments of semantic relatedness and also the neural organization of ITC.

Correspondences between Fine Grain Structure in Semantic Similarity and ITC

Each of our measures of semantic similarity exhibited a qualitative three/four category structure to varying degrees, whereas ITC exhibited a two-category structure based on animacy. This global mismatch might explain the relatively poor performance of the WordNet measures. To study the fine grain structure of the representations, we examined the correspondence between ITC's representation and our semantic measures for "within" and "between" category comparisons separately (see Methods). By parceling the data in this way, we can examine correspondences between ITC and semantic measures in terms of (a) fine-grained substructure within categories (e.g., within the category human are "woman" and "chef") and (b) associations between objects in different categories (e.g., across the categories human and man-made objects are "woman" and "umbrella"). Note the mismatches associated with different category structure in ITC and in the semantic measures will be reflected in relatively low correspondence for the between category comparisons.

The results of the analysis are shown in Figure 3. For the analysis of within category structure, we found a correspondence between ITC's representation and explicit ratings and COALS, showing that ITC represents fine-grained relationships between exemplars within a category in a way comparable to these measures. Notably, LSA performed marginally well (as did WordNet VECTOR) but did not reach statistical significance ($p > .05$). In light of LSA's good performance overall (see above), the failure to reach significance for the within-category comparisons might simply result from a smaller parceled data set. The

PATH and GLOSS WordNet measures, in contrast, had virtually no relationship with ITC's within category structure. Consistent with the observed global mismatch, we find no correspondence between ITC's representation and any of the semantic measures for the analysis of between category structure, indicating that associations between categories might be encoded elsewhere in the brain, possibly in conjunction with ITC.

DISCUSSION

In this study, we examined whether the organization of semantic knowledge is reflected in the brain's representation of information in ITC, an area traditionally thought of as a visual area (Logothetis & Sheinberg, 1996; Gross, 1992). We studied the relationship between six measures of semantic relatedness (explicit ratings, three lexicographic measures from WordNet, and two emergent measures from distributional analyses) and neural representations in both visual cortex and ITC. We found that the semantic relationships between the labels ascribed to visual objects are reflected in the brain's representations of these objects in ITC. More specifically, we found a correspondence between the geometry of ITC representations and semantic relationships, as expressed in explicit behavioral ratings and in distributional measures of semantic similarity, and that this connection largely could be ascribed to correspondences in the encoding of within category relationships.

In our study, we found a coupling between the semantic relationships and the brain's representation of information in ITC. This coupling is not perfect: We observed a coarse mismatch in the broad categorical boundaries of the representations in ITC (animate/inanimate) and the measures of semantic similarity (human, animal, natural, and man-made objects), which was reflected in a poor correspondence for between-category structure. One possibility is that this broad categorical structure draws primarily on nonperceptual relationships. For example, humans are relatively highly related to all objects in the explicit judgments data (note relatively high similarities in the leftmost and topmost columns of the explicit similarity judgments DSM), perhaps reflecting associative knowledge that is distinct from conceptual similarity per se (e.g., humans use tools, eat food, etc.). In any case, this mismatch suggests that the brain's underlying neural representation of semantic relationships involves more than ITC. Mitchell et al. (2008) have argued that semantic meaning is encoded in a network of brain activity, consisting of frontal areas and sensory motor areas in addition to ITC. Our findings, in the context of their proposal, suggest that ITC is an important node in this network, as activity in ITC (studied in isolation from the network) reflects the topology of semantic relationships within categories.

In contrast to our finding that semantic structure, as expressed in word use and explicit judgments, is represented to some extent at the level of ITC, we found no

correspondence between any measure of semantic similarity and activity in primary visual cortex. This pattern indicates that semantic relationships may be an organizational principle in ITC and that the refinement of visual inputs into linguistically and conceptually relevant organization begins in the ventral visual pathway. Although this suggests that semantic knowledge could be represented in ITC, our findings do not necessarily mean that ITC is playing an explicit role in linguistic semantics. Developmental studies have suggested a moderated relationship between perceptual categories, conceptual categories, and language. Early theories of development pointed out these capacities have different developmental trajectories, implying separate interrelated systems (Piaget, 1952). Contemporary theories view the development of perceptual categories as the building blocks of conceptual categories and language (Mandler, 2004; Karmiloff-Smith, 1992), explicating the dependencies between perception, concepts, and language. The observed correspondence between semantic relationships and ITC's representation of information may therefore reflect the perceptual origins of conceptual categories and language.

Although ITC certainly represents perceptual information, our findings are also congruent with the idea that representations in ITC are not purely perceptual. Nonvisual tasks based on conceptual knowledge modulate ITC activity (for a review, see Martin, 2007), and Simanova et al. (2012) recently showed that, irrespective of sensory modality, the categories of animals and tools could be decoded from ITC activity. These findings suggest that ITC's representations may be more conceptual than sensory, fitting with our finding that ITC's representation of visual objects was associated with representations based on the word labels associated with those objects, which abstract away from sensory information. Furthermore, by studying the relationships between a large number of stimuli (summarized by the DSMs), we show a level of correspondence between semantic relationships and ITC's representation that is far more sophisticated than the coarse animal/tool categorical distinction demonstrated by Simanova et al. (2012). The observation that the complex geometry of semantic relationships among words matches with the representational geometry of ITC strengthens the argument that conceptual knowledge is represented, at least to some extent, in ITC. Moreover, our findings showing compatibility at a fine grain for within category relationships further specify the role that ITC might play in representing conceptual knowledge.

It is clear that the topology of ITC does not completely capture amodal semantic organization given that the overall correlations we observe with the language-based measures are relatively low. The difference between within-category and between-category relationships suggests that this reflects, at least partially, a mismatch between high-level categorical structure in ITC and the semantic measures (as correlations between ITC and semantic measures are higher within categories). An

additional possibility is that ITC representations primarily capture visually based aspects of conceptual knowledge whereas the semantic relationships measured by LSA, COALS, and the behavioral judgments incorporate both visually based and nonvisual aspects of meaning. Although this seemingly contrasts with the findings of Simanova et al. (2012), discussed above, their data are also consistent with sensory-based representations in ITC if the relevant visual features are coactivated from activation in other modalities (e.g., the auditory word “red” might activate simulations in color areas). At the extreme, it seems implausible that visual areas represent the semantics of nonvisual abstract concepts like “love” or “morality,” although visual areas might still represent visually based associations related to these nonvisual concepts. By this account, ITC plays an important role primarily in the representation of visually based aspects of conceptual knowledge as reflected in language use, contributing to larger distributed multi-modal representations of semantic meaning.

A related account is that both linguistic use and the representation of information ITC cortex reflect a similarity that exists “in the world”—that is, our environment separately drives ITC organization and the organization of semantic knowledge. Conceptual knowledge likely draws on information from multiple modalities (cf. Mahon & Caramazza, 2011), presumably requiring some degree of commonality in the formatting of information across systems. Common pressure from the organizational principles of the environment may naturally lead to this “common code,” thereby allowing for efficient communication and integration across sensory and cognitive neural systems (cf. Prinz, 1997). Indeed, some recent fMRI studies on congenitally blind participants demonstrate innate conceptual biases in the organization of object knowledge independent of sensory input modality and experience (Striem-Amit, Cohen, Dehaene, & Amedi, 2012; Reich, Szwed, Cohen, & Amedi, 2011; Mahon, Anzellotti, Schwarzbach, Zampini, & Caramazza, 2009), indicating that multimodal pressures are influencing the organizational principles of the ventral visual pathway.

One limitation with these data is that they are derived from a single language (English). This is a potential limitation (of this and of many other studies; cf. Henrich, Heine, & Norenzayan, 2010) because languages differ in how they categorize entities and so speakers of different languages may show corresponding differences in how they map representations in ITC onto semantic representations. Such cross-linguistic differences can be quite substantial, including different ways of grouping important conceptual categories like body parts (Brown, 2011; Majid, 2010) and natural objects (e.g., Levinson, 1996). It would be valuable for future work to investigate the relationship between the topology of ITC and of language-based measures in speakers of other languages as this would inform the extent of interaction between the representational structures of ITC and language. Some necessary ground-

work for this goal comes from work building versions of WordNet and language corpora in other languages and in multilingual contexts (Vossen, 1998) and work demonstrating the utility of relatedness measures in those contexts (Hassan, Banea, & Mihalcea, 2012; Mohammad, Gurevych, Hirst, & Zesch, 2007; Katz & Goldsmith-Pinkham, 1998). If these cross-linguistic differences are in fact reflected in the topology of ITC, this would support a strong relationship between visual processing and language-based semantic structure.

The Mismatch of WordNet with ITC

The structure of semantic knowledge in WordNet was not reflected in ITC’s representation. Our qualitative analysis of the DSMs for WordNet and explicit behavioral ratings found a three/four category structure in the semantic relatedness data (human, animal, natural objects, and man-made objects). LSA and COALS similarly exhibited this structure (see DSMs in Figure 2). The presence of this structure in LSA, COALS, and behavioral ratings, all of which correlated with ITC structure, rules out the possibility that our null results for the WordNet measures were because of a coarse hierarchical mismatch. This is supported by the outcome of our fine-grained analysis for within and between category structure. Even after this coarse mismatch was parceled out (by looking separately at organizational similarity between and within categories), none of the WordNet measures captured the fine grained within category structure. The lack of a correspondence between WordNet and ITC thus is more likely attributed to more subtle differences between WordNet’s topology and ITC’s representation. These subtle differences could be ascribed to WordNet’s (imposed) structure and word definitions (glosses), which might fail to fully capture the structure of our mental conceptual representations. Volumes of text have been written to describe concepts that are captured by a single word (e.g., a search for nonfiction books using the keyword “monkey” in the University of Maryland’s library produced nearly 1400 results), and short definitions from different sources are often quite different. WordNet’s simple structure and short definitions may just be too simplistic to capture the full range of word meaning.

Relatedly, the WordNet hierarchy may be too sparsely populated (i.e., too few links between word synsets) and qualitative (e.g., the distance between “run” and “jog” is assumed to be the same as between “run” and “move”) to adequately predict cognitive representations of word meaning (Boyd-Graber & Fellbaum, 2006). In contrast, LSA and COALS, which incorporate volumes of text to derive the relations among words, may unsurprisingly better reflect both explicit similarity ratings and the organization of information in human ITC. An alternative explanation for WordNet’s lack of correspondence with ITC is that the structure of semantic meaning as represented in WordNet might be reflected in other areas of the brain.

Our study is limited in addressing this possibility as we based our research on an existing data set that only includes data from the primary visual cortex and ITC. Future research employing a similar approach could address this possibility more directly.

Future Directions for Modeling of Semantic Knowledge

The development of WordNet measures and distributional measures like LSA has been critical to the advancement of our understanding of how the brain encodes semantic knowledge. Some work takes these measures (especially distributional measures like LSA and COALS) as a veridical account of semantic representation in the human brain, at least implicitly (e.g., Landauer & Dumais, 1997); however it is possible that these measures relate well to semantic knowledge without corresponding directly to the neural representation of semantic or conceptual knowledge (Glenberg & Mehta, 2008). Nevertheless, such metrics can capture an important part of the “material” available for learning semantic structure (e.g., Andrews, Vigliocco, & Vinson, 2009), and our data suggest that measures like LSA do, in fact, reflect important aspects of how the human brain represents conceptual information.

The highly active and competitive field of constructing models to describe semantic relations among words most often uses behavioral measures as benchmark (e.g., Kievit-Kylar & Jones, 2012; Riordan & Jones, 2010). Many studies rating semantic measures using human judgments use pre-compiled data sets (see Finkelstein et al., 2002; Miller & Charles, 1991; Rubenstein & Goodenough, 1965) and typically find that distributional measures like COALS and LSA outperform WordNet-based measures (Waltinger & Mehler, 2009; Boyd-Graber & Fellbaum, 2006; Rohde et al., 2005). However, some WordNet algorithms that we did not implement in the current study may correlate more strongly with the behavioral data (e.g., Budanitsky & Hirst, 2005; Lapata & Barzilay, 2005; Jarmasz & Szpakowicz, 2003). A more nuanced sense disambiguation heuristic may also improve WordNet performance (see Methods). In addition, there are likely several other types of hierarchical and distributional models worth investigating (e.g., Panchenko, 2012; Mohammad & Hirst, 2006), including models based on semantic feature norms (McRae, Cree, Seidenberg, & McNorgan, 2005; Cree & McRae, 2003).

We would argue that this field of research could be additionally served by employing “brain-based” benchmarks, both those targeting specific areas, as in this study, and those using whole-brain analysis approaches (e.g., Huth et al., 2012; Mitchell et al., 2008). It may be the case that hierarchical and distributional measures of semantic similarity are only measuring specific aspects of linguistic knowledge (see Glenberg & Mehta, 2008; Maki & Buchanan, 2008); brain-based measures potentially

could serve to reveal shortcomings of specific models and provide guidance for their refinement. In addition, there is the potential of improving these measures by integrating feature sets and distributional data (cf. Andrews et al., 2009), representing word meanings probabilistically (see Griffiths, Steyvers, & Tenenbaum, 2007; Blei, Ng, & Jordan, 2003) and using other types of text corpora such as the Google Books corpus or Wikipedia (Ferrara & Tasso, 2013; Michel et al., 2011; Gabrilovich & Markovitch, 2007). By making use of existing data sets, as in this study, this research could be performed at a relatively low cost.

APPENDIX 1

Word Labels for the Stimuli

<i>Human</i>	<i>Animal</i>	<i>Natural Objects</i>		<i>Man-made Objects</i>	
hand	armadillo	carrots	garlic	city	umbrella
ear	camel	grapes	path	bottle	phone
chef	snake	potato	pineapple	lightbulb	stove
hair	wolf	tree	pear	sign	
dancer	monkey	pepper	waterfall	cassette	
woman	ostrich	lettuce		church	
eye	zebra	kiwi		flag	
man	elephant	cucumber		key	
pointing	sheep	leaf		pliers	
fist	frog	apple		arch	
child	cow	radish		door	
	goat	eggplant		hammer	
	dog	lake		chair	
	alligator	pine cone		gun	
	giraffe	banana		house	
	lion	tomato		dome	

Reprint requests should be sent to Thomas A. Carlson, Department of Cognitive Sciences, Centre for Cognition and Its Disorders, Macquarie University, Sydney, NSW 2109, Australia, or via e-mail: thomas.carlson@mq.edu.au.

REFERENCES

- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, 116, 463–498.
- Banerjee, S., & Pedersen, T. (2002). An adapted Lesk algorithm for word sense disambiguation using WordNet. In *Proceedings of the Third International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2002)* (pp. 136–145), Mexico City, Mexico.

- Banerjee, S., & Pedersen, T. (2003). Extended gloss overlaps as a measure of semantic relatedness. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence* (pp. 805–810), Acapulco, Mexico.
- Basso, A., Capitani, E., & Laiacona, M. (1988). Progressive language impairment without dementia: A case with isolated category specific semantic defect. *Journal of Neurology, Neurosurgery, and Psychiatry*, *51*, 1201–1207.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3*, 993–1022.
- Boyd-Graber, J., & Fellbaum, C. (2006). Adding dense, weighted connections to wordnet. In *Proceedings of the 3rd Global WordNet Meeting* (pp. 29–35), Jeju-do, Korea.
- Brown, C. H. (2011). Hand and arm. In M. S. Dryer & M. Haspelmath (Eds.), *The world atlas of language structures* (vol. 100, pp. 522–525). Munich: Max Planck Digital Library. Retrieved from wals.info/chapter/129.
- Budanitsky, A., & Hirst, G. (2005). Evaluating WordNet-based measures of lexical semantic relatedness. *Computational Linguistics*, *32*, 13–47.
- Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*, 213–261.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain the animate-inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., et al. (2012). The representation of biological classes in the human brain. *The Journal of Neuroscience*, *32*, 2608–2618.
- Cree, G. S., & McRae, K. (2003). 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*, *132*, 163–201.
- Dahl, Ö. (2008). Animacy and egophoricity: Grammar, ontology and phylogeny. *Lingua*, *118*, 141–150.
- Dahl, Ö., & Fraurud, K. (1996). Animacy in grammar and discourse. In T. Fretheim & J. K. Gundel (Eds.), *Reference and reference accessibility* (pp. 47–64). Amsterdam: Benjamins.
- Davis, M. H., & Johnsrude, I. S. (2003). Hierarchical processing in spoken language comprehension. *The Journal of Neuroscience*, *23*, 3423–3431.
- Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*, 2470–2473.
- Fellbaum, C. (1998). *WordNet: An electronic lexical database*. Cambridge, MA: MIT Press.
- Ferrara, F., & Tasso, C. (2013). Evaluating the results of methods for computing semantic relatedness. *Computational Linguistics and Intelligence Text Processing*, *7816*, 446–458.
- Ferreira, F., & Tanenhaus, M. K. (2007). Introduction to the special issue on language–vision interactions. *Journal of Memory and Language*, *57*, 455–459.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., et al. (2002). Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, *20*, 116–131.
- Gabrilovich, E., & Markovitch, S. (2007). Computing semantic relatedness using Wikipedia-based explicit semantic analysis. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 1606–1611).
- Giraud, A. L. (2004). Contributions of sensory input, auditory search and verbal comprehension to cortical activity during speech processing. *Cerebral Cortex*, *14*, 247–255.
- Glenberg, A. M., & Mehta, S. (2008). Constraint on covariation: It's not meaning. *Rivista di Linguistica*, *20*, 237–262.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*, 211–244.
- Gross, C. G. (1992). Representation of visual stimuli in inferior temporal cortex. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *335*, 3–10.
- Hallgren, K. A. (2012). Computing interrater reliability for observational data: An overview and tutorial. *Tutorials in Quantitative Methods for Psychology*, *8*, 23–34.
- Hassan, S., Banea, C., & Mihalcea, R. (2012). Measuring semantic relatedness using multilingual representations. In *First Joint Conference on Lexical and Computational Semantics* (pp. 20–29), Montreal, Canada.
- Hawker, T., & Honnibal, M. (2006). Improved default sense selection for word sense disambiguation. In *Proceedings of the 2006 Australasian Language Technology Workshop (ALTW2006)* (pp. 11–17), Sydney, Australia.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *The Behavioral and Brain Sciences*, *33*, 61–83; discussion 83–135.
- Huth, A. H., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, *76*, 1210–1224.
- Jarmasz, M., & Szpakowicz, S. (2003). Roget's thesaurus and semantic similarity. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2003)* (pp. 212–219), Borovets, Bulgaria.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, *17*, 4302–4311.
- Kamiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT Press.
- Katz, P., & Goldsmith-Pinkham, P. (1998). *Word sense disambiguation using latent semantic analysis*. Retrieved from www.wsrn.sccc.swarthmore.edu/users/07/pkatz1/cs65f06-final.pdf.
- Kiani, R., & Esteky, H. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of Physiology*, *4296*–4309.
- Kievit-Kylar, B., & Jones, M. N. (2012). Visualizing multiple word similarity measures. *Behavior Research Methods*, *44*, 656–674.
- Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, *74*, 1114–1124.
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., et al. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*, 1126–1141.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.
- Lapata, M., & Barzilay, R. (2005). Automatic evaluation of text coherence: Models and representations. In *19th International Joint Conference on Artificial Intelligence* (pp. 1085–1090), Edinburgh. Retrieved from ijcai.org/Past Proceedings/IJCAI-05/PDF/0505.pdf.
- Levinson, S. C. (1996). Relativity in spatial conception and description. In J. J. Gumperz & S. C. Levinson (Eds.),

- Rethinking linguistic relativity* (pp. 177–202). Cambridge: Cambridge University Press.
- Logothetis, N., & Sheinberg, D. (1996). Visual object recognition. *Annual Review of Neuroscience*, *19*, 577–621.
- Mahon, B. Z., Anzellotti, S., Schwarzbach, J., Zampini, M., & Caramazza, A. (2009). Category-specific organization in the human brain does not require visual experience. *Neuron*, *63*, 397–405.
- Mahon, B. Z., & Caramazza, A. (2011). What drives the organization of object knowledge in the brain? *Trends in Cognitive Sciences*, *15*, 97–103.
- Majid, A. (2010). Words for parts of the body. In B. C. Malt & P. Wolff (Eds.), *Words and the mind: How words capture human experience* (pp. 58–71). Cambridge: Cambridge University Press.
- Maki, W. S., & Buchanan, E. (2008). Latent structure in measures of associative, semantic, and thematic knowledge. *Psychonomic Bulletin & Review*, *15*, 598–603.
- Mandler, J. M. (2004). Thought before language. *Trends in Cognitive Sciences*, *8*, 508–513.
- Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, *58*, 25–45.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, *37*, 547–559.
- Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., et al. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, *331*, 176–182.
- Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, *38*, 39–41.
- Miller, G. A., & Charles, W. G. (1991). Contextual correlates of semantic similarity. *Language and Cognitive Processes*, *6*, 1–28.
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., et al. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, *320*, 1191–1195.
- Mohammad, S., Gurevych, I., Hirst, G., & Zesch, T. (2007). Cross-lingual distributional profiles of concepts for measuring semantic distance. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP/CoNLL-2007)*, Prague, Czech Republic.
- Mohammad, S., & Hirst, G. (2006). *Distributional measures as proxies for semantic distance: A survey*. Retrieved from www.umiacs.umd.edu/~saif/WebDocs/distributionalmeasures.pdf.
- Moldovan, D., & Novischi, A. (2004). Word sense disambiguation of WordNet glosses. *Computer Speech & Language*, *18*, 301–317.
- Panchenko, A. (2012). A study of heterogeneous similarity measures for semantic relation extraction. In *Proceedings of 14e Rencontres des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (JEP-TALN-RECITAL), 2012* (vol. 3, pp. 29–42), Grenoble, France.
- Patwardhan, S., & Pedersen, T. (2006). Using WordNet-based context vectors to estimate the semantic relatedness of concepts. In *Proceedings of the EACL 2006 Workshop: Making Sense of Sense—Bringing Computational Linguistics and Psycholinguistics Together* (pp. 1–8), Trento, Italy. Retrieved from acl.ldc.upenn.edu/eacl2006/ws09_makingsense.pdf#page=7.
- Pedersen, T., Patwardhan, S., & Michelizzi, J. (2004). WordNet::Similarity: Measuring the relatedness of concepts. In *Proceedings of the 5th Annual Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL-04)* (pp. 38–41), Boston, Massachusetts. Retrieved from dl.acm.org/citation.cfm?id=1614025.1614037.
- Piaget, J. (1952). *The origin of intelligence in the child*. New York: International University Press.
- Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, *9*, 129–154.
- Reddy, L., & Kanwisher, N. (2006). Coding of visual objects in the ventral stream. *Current Opinion in Neurobiology*, *16*, 408–414.
- Reich, L., Szwed, M., Cohen, L., & Amedi, A. (2011). A ventral visual stream reading center independent of visual experience. *Current Biology: CB*, *21*, 363–368.
- Riordan, B., & Jones, M. N. (2010). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science*, *3*, 303–345.
- Rodd, J. M., Davis, M. H., & Johnsrude, I. S. (2005). The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. *Cerebral Cortex*, *15*, 1261–1269.
- Rohde, D. L. T., Gonnerman, L. M., & Plaut, D. C. (2005). *An improved model of semantic similarity based on lexical co-occurrence*. Unpublished manuscript. Retrieved from tedlab.mit.edu/~dr/Papers/RohdeGonnermanPlaut-COALS.pdf.
- Rubenstein, H., & Goodenough, J. B. (1965). Contextual correlates of synonymy. *Communications of the ACM*, *8*, 627–633. Retrieved from dl.acm.org/citation.cfm?id=365657.
- Silveri, M. C., & Gainottib, G. (1988). Interaction between vision and language in category-specific semantic impairment. *Cognitive Neuropsychology*, *5*, 677–709.
- Simanova, I., Hagoort, P., Oostenveld, R., & Van Gerven, M. A. J. (2012). Modality-independent decoding of semantic information from the human brain. *Cerebral Cortex*. doi:10.1093/cercor/bhs324.
- Striem-Amit, E., Cohen, L., Dehaene, S., & Amedi, A. (2012). Reading with sounds: Sensory substitution selectively activates the visual word form area in the blind. *Neuron*, *76*, 640–652.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.
- Vossen, P. (1998). *EuroWordNet: A multilingual database with lexical semantic networks. Psychoanalytic review* (vol. 100). Dordrecht, The Netherlands: Kluwer Academic Publishers. doi:10.1521/prev.2013.100.3.513.
- Waltinger, U., & Mehler, A. (2009). Social semantics and its evaluation by means of semantic relatedness and open topic models. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology* (pp. 42–49). IEEE. doi:10.1109/WI-IAT.2009.14.