Similarity judgment within and across categories: A comprehensive model comparison

Russell Richie* & Sudeep Bhatia

University of Pennsylvania, USA

April 2, 2021

Author Note

* Correspondence regarding this article should be addressed to Russell Richie, 3720 Walnut St., Philadelphia, PA 19104, drichie@sas.upenn.edu. Funding was received from the National Science Foundation grant SES-1847794 and the Alfred P. Sloan Foundation.

Abstract

Similarity is one of the most important relations humans perceive, arguably subserving category learning and categorization, generalization and discrimination, judgment and decision making, and other cognitive functions. Researchers have proposed a wide range of representations and metrics that could be at play in similarity judgment, yet have not comprehensively compared the power of these representations and metrics for predicting similarity within and across different semantic categories. We performed such a comparison by pairing nine prominent vector semantic representations with seven established similarity metrics that could operate on these representations, as well as supervised methods for dimensional weighting in the similarity function. This approach yields a factorial model structure with 126 distinct representation-metric pairs, which we tested on a novel dataset of similarity judgments between pairs of co-hyponymic words in eight categories. We found that cosine similarity and Pearson correlation were the overall best performing unweighted similarity functions, and that word vectors derived from free association norms often outperformed word vectors derived from text (including those specialized for similarity). Importantly, models that used human similarity judgments to learn category-specific weights on dimensions yielded substantially better predictions than all unweighted approaches across all types of similarity functions and representations, although dimension weights did not generalize well across semantic categories, suggesting strong category context effects in similarity judgment. We discuss implications of these results for cognitive modeling and natural language processing, as well as for theories of the representations and metrics involved in similarity.

Keywords: similarity; vector semantics; association norms; computational models; context

1. Introduction

Much of cognition and behavior relies on some notion of similarity. The probability of generalizing properties from one stimulus to another is proportional to how similar they are (Shepard, 1987, 2004), and categorization of a stimulus into one class instead of another correlates with the relative similarity of the stimulus to the two classes (Nosofsky, 1984, 1986). Additionally, various complex judgments and decisions, including multiattribute choice, probability judgment, and social judgment, all rely on similarity to some extent (Tversky, 1969; Kahneman & Tversky, 1972; Tversky & Kahneman, 1983). Similarity may also play a fundamental role in memory search, with retrieved items cuing the subsequent retrieval of similar items (Bousfield & Sedgewick, 1944; Gruenewald & Lockhead, 1980; Howard & Kahana, 2002; Romney et al., 1993; Hills et al., 2012). Relatedly, many types of queries or recommendations, to ourselves, other humans, or artificial systems, also require returning items similar to a probe, as in “I want a career change. What’s a job similar to a lawyer?”, or “Siri, I’m hungry. What’s a fruit like an apple?”. In each case, the query recipient must find professions and fruits similar to ‘lawyer’ and ‘apple’, respectively.

Modeling and predicting similarity is thus of great importance to cognitive scientists. Under one common conception in cognitive science, constructing a model of similarity judgment requires a specification of both the information that is represented, and how that information is operated upon to produce a judgment or perception of similarity. Accordingly, cognitive science has produced many candidate theories of representations and metrics underlying similarity (for reviews, see Goldstone and Son, 2004; Hahn, 2014; Jones et al., 2015). Many of these metrics involve the explicit or implicit use of distance functions applied to spatial or feature

representations (Shepard & Arabie, 1979; Nosofsky, 1986; Shepard, 1987). More recently, semantic vector models of words and phrases derived from text (Mikolov et al., 2013; Pennington et al., 2014) and free association norms (De Deyne et al., 2019), have been used as proxies for the item representations at play in similarity judgment (largely due to their ability to accurately scale to thousands or even millions of naturally occurring items over which similarity metrics could operate). As different types of representations (e.g., different semantic vector models) could, in principle, be combined with different types of metrics (e.g., different similarity functions operating over semantic spaces), it is crucial to systematically compare numerous possible representation-metric pairs for predicting similarity judgment.

There have been some attempts at comparing different similarity models, both in cognitive science (Landauer & Dumais, 1997; Burgess & Lund, 2000; Rohde et al., 2006; Bullinaria & Levy, 2007, 2012, Mandler et al., 2017; and Pereira et al., 2016) and computational linguistics (e.g., Hill et al., 2014, 2015; Wieting et al., 2015; Gerz et al., 2016; Ponti et al., 2020). However, most of this work has used benchmark datasets (e.g. the TOEFL dataset and SimLex-999; Landauer & Dumais, 1997; Hill et al., 2015) that sample pairs of words from the entire lexicon and include a large number of rather unrelated pairs of words (for a similar point, see De Deyne et al., 2021). For example, in SimLex-999, pairs include ‘wife-straw’ and ‘ankle-window’. Thus, the high performance of existing models in these benchmark datasets could be largely due to the ability of the models to differentiate between general categories (such as objects, human relations, and body parts) rather than their ability to predict similarity structures within categories. That is, these models might struggle when differentiating the level of similarity between co-hyponyms – words belonging to the same hypernym or category – like

‘wife-husband’ and ‘wife-uncle’ (which share the category ‘family relations’), or ‘ankle-knee’ and ‘ankle-face’ (which share the category ‘body parts’). Of course, cognitive scientists are often interested in similarity judgments for co-hyponymic items within a specific category (e.g., Russell, 1980; Nosofsky, 1984, 1986; Hills et al., 2012; Hout et al., 2013; Wnuk & Majid, 2014), which may be more challenging to predict, as it requires making more fine-grained distinctions among items. It could also be the case that a given similarity model is good at making predictions for items within some categories, but not others.

Additionally, most previous work evaluating models of similarity judgment has examined a relatively small set of metrics (e.g. three distance functions) that could be operating on the representations. In fact, the focus of this work has largely been on evaluating different ways of generating semantic representations (e.g. different text processing assumptions and learning algorithms for word vector models – see Bullinaria & Levy, 2007, 2012, Mandera et al., 2017), and not on comparing pairs of distinct representation and metric combinations. Yet metrics and representations could interact, yielding emergent properties of similarity models not explored in prior work. Further, there has been little investigation of the possibility that *weighted* functions of similarity over vector representations, with weights obtained using a sample of participant judgments of similarity, make better predictions of similarity than unweighted functions (but see Peterson et al., 2018 for applications of this idea to similarity judgments involving images, and Chen et al., 2017 and Lu et al., 2019 for relational judgments between words). For this reason, there has also been little work attempting to understand how different distance functions (including functions with different sets of weights) could be operating in different categories.

The goal of the present work, therefore, is to evaluate different representation-metric pairs' ability to predict similarity judgments among co-hyponyms, in a full representation-by-metric factorial design. For weighted similarity functions, we consider functions learned from participant judgments on a per-category basis, as well as those learned from participant judgments of a combination of categories all at once. We focus, in particular, on similarity models that make predictions for the tens of thousands of words that could be evaluated in similarity judgment. The generality of our tests allows us to not only understand if some representations or metrics generally make better predictions of similarity, but also determine if some representation-metric combinations are especially predictive of similarity, and how this varies across different similarity judgment domains, categories, or contexts. Further, by using participant judgments to determine the dimensional weights that best predict similarity, we can calibrate similarity models to improve predictive accuracy and better understand how similarity metrics might vary across semantic categories or contexts.

Overall, our tests shed light on the representations and metrics that are likely to underlie the perception of similarity, and yield practical implications in terms of convenient but accurate approximation to similarity data for various cognitive science and natural language processing applications. In the sections that follow, we describe in more detail the representations and similarity functions that we consider.

1.1. Representations

1.1.1 Text-based representations

There are at least two possible sources for rich, general-purpose representations of the thousands or even millions of words whose similarity we might perceive. First, high-

dimensional, real-valued vector representations for millions of words and phrases can be constructed based on lexical co-occurrence statistics in large natural language corpora. While there are numerous ways of constructing such representation, we focus in this work on word2vec (Mikolov et al., 2013), fastText (Mikolov et al., 2018), and GloVe (Pennington et al., 2014). These are powerful new methods that, unlike older techniques in cognitive science (e.g. Griffiths et al., 2007; Jones & Mewhort, 2007; Landauer & Dumais, 1997), have been applied to billions of words of natural language data to generate rich representations for a large number of words. There are two complementary techniques used in word2vec: skip-gram with negative sampling (SGNS) and continuous bag-of-words (CBOW). Skip-gram relies on a multilayer feedforward neural network that slides over windows of text in a corpus and attempts to predict the words in the periphery of the window, given the word in the center of the window. CBOW attempts to do the inverse (predict the word in the center given the periphery). The rows of the trained weight matrix from the input layer to the hidden layer are the word vector representations typically used for SGNS, while the weight matrix from the hidden layer to the output layer are the representations typically used for CBOW. The fastText algorithm also uses CBOW, but also exploits subword information so that vectors can be generated for new words morphologically related to words in the training corpus (e.g., if the corpus contains ‘veterinarian’ but not ‘veterinarians’, fastText will still be able to generate a vector for the latter which will be closely related to the former). GloVe first constructs a global word-word co-occurrence matrix whose entries encode the frequency with which words co-occur with one another in a given corpus. Vectors are then learned such that their dot product equals the logarithm of words’ probability of co-occurrence.

These representations are useful for modeling semantic priming, free association, list recall, and a host of other phenomena (see Bhatia et al., 2019; Jones et al., 2015; Lenci, 2018; Mander et al., 2017 for reviews). Most relevant, standard word vector models like word2vec and GloVe can model a range of semantic relatedness judgments (e.g., doctor-hospital or bad-good), usually by correlating numerical ratings of two words' relatedness with the two words' vector similarity. However, these text-based models struggle to account for judgments when participants are asked to rate true similarity (e.g., doctor-surgeon) instead of generic relatedness (e.g., doctor-hospital are related but not similar) or to recognize that antonyms (e.g., good-bad) are dissimilar, as in the popular similarity benchmark SimLex-999. In response, approaches for 'specializing' these vector spaces to respect true similarity relations have been developed. For example, Paragram-SL999 uses a large set of lexical 'paraphrase pairs' (basically, synonyms or similar word pairs) to adjust pre-trained vectors such that vectors for paraphrase pairs are moved closer together (Wieting et al., 2015). While cosine similarity between standard word2vec vectors correlates with SimLex-999 judgments at $r = .37$, Paragram-SL999 vectors achieve a correlation of $r = .68$. This approach is successful on SimLex-999, and improves accuracy in a sentiment analysis task, but the downside of this approach is that only the vectors for the ~100,000 word pairs in the paraphrase dataset are adjusted. The adversarial post-specialization method attempts to solve this problem by using synonym and antonym pairs to learn a global specialization function that can be applied to all word vectors, i.e., move all words (even those not in the synonym/antonym pairs) so that they better respect similarity relations (Ponti et al., 2020). Cosine similarities between GloVe vectors transformed with this approach correlate with human judgments in SimLex-999 at $r = .78$, which is also the average correlation between one

human SimLex-999 rater and the average of the rest, thus suggesting that this approach has possibly reached the ceiling of inter-rater reliability. Vectors specialized with this approach also improve performance over their unspecialized counterparts in natural language processing tasks like lexical simplification (replacing complex words with simpler synonyms in a sentential context) and dialogue state tracking (estimating the beliefs of a human question-asker at every dialogue turn; Ponti et al., 2020; Mrkšić et al., 2017).

How should we interpret the (differences among) text-based representations in potentially psychologically relevant ways? We point interested readers to other more detailed treatments of these models from a cognitive perspective (e.g., Günther et al., 2019; Mander et al., 2017), but briefly, one common way to differentiate among text-based models is in terms of counting vs prediction methods. GloVe is considered the former type, as it involves explicitly counting word-word co-occurrences in a large matrix. As waiting to build representations until after an entire corpus has been collected seems cognitively implausible, counting models are viewed as computational- and not algorithmic-level theories (Marr, 1982). Word2vec and fastText, however, are prediction models, as they involve neural networks that take text in word-by-word, and attempt to predict words in context. Unlike counting models, prediction models might be plausible algorithmic-level theories, as prediction models are akin to classic psychological models of associative learning like Rescorla-Wagner (Mander et al., 2017). It is important to note, however, that there is some provable equivalence between prediction and counting methods. For example, word2vec is equivalent to factorization of a word-context matrix, whose cells are the pointwise mutual information of the respective word and context pairs, shifted by a global constant (Levy & Goldberg, 2014; see also Landauer & Dumais, 1997,

for an interpretation of Latent Semantic Analysis in terms of a 3-layered neural network).

Therefore, any counting method – while not an algorithmically-plausible model itself – may correspond to an algorithmically-plausible model of learning. In any case, because of the provable equivalence between counting and prediction models, and because prior work shows little difference between these models when hyperparameters are carefully tuned (Levy et al., 2015), we do not expect large differences between, say, fastText and GloVe, when they are trained on comparably sized corpora. Instead, we are more interested in comparing unspecialized embeddings to specialized embeddings (and to the association norm-derived vectors we describe in the next section).

Like counting models, it's also dubious that the post-specialization techniques, at least as described above, embody a psychologically plausible theory of learning, since they pre-train word embeddings in one stage, and then adjust them in a second stage with synonym/antonym constraints. However, it is also possible to incorporate such constraints during a single stage of learning from text (Ponti et al., 2020). If such a single-stage approach is instantiated in a prediction model (cf. count model), then a plausible psychological model of learning from textual co-occurrence and knowledge of lexical relations may be possible.

Of course, all our models were trained on billions of words, which is more than any human consumes in a lifetime, so our models do not constitute complete models of *how* semantic representations are obtained, or what they are learned *from*. Instead, we view them as computational-level theories of *what* semantic representations *are learned* (from unknown learning environments and procedures). It could be that almost all of one's concepts for 'cat', 'dog', or 'bird' for example, come from one's experience with cats, dogs, and birds in the world.

Nonetheless, the good performance of text-based representations in various behavioral tasks (again, see Bhatia et al., 2019; Jones et al., 2015; Lenci, 2018; Mandera et al., 2017) suggest that there is at least some isomorphism between what might be learned from those real-world experiences, and what can be learned from text.

1.1.2. Association-based representations

The second general-purpose representation comes from free association norms, where participants are given a cue word and asked to respond with whatever words come to mind (De Deyne, et al., 2019). For example, given the cue ‘cow’, a participant might respond with ‘domesticated’ or ‘herbivorous’ as they could in a feature norm study, but might also respond with ‘farm’ or ‘horse’. These latter responses are of course not features or properties of cows, but rather words or concepts with which cows are semantically associated (because cows often *live on farms*, *with horses*). Such responses can be aggregated into a cue-response matrix by counting how often particular words or phrases (columns) were generated in response to particular cue words (rows). There are various ways to process this information and use it to make predictions of similarity judgments (or other behavioral data). One moderately successful approach is to transform each word-response cell to its positive pointwise mutual information, which effectively downweights responses that are given often regardless of the cue:

$$\text{PPMI}(r \mid c) = \max\left(0, \log_2\left(\frac{p(r \mid c)N}{\sum_i p(r \mid c_i)}\right)\right)$$

where N is the total number of cues, and $p(r/c)$ is the probability of response r given cue c .

However, De Deyne et al. found that the best procedure entailed transforming the matrix such that a cue-response cell represents, approximately, the activation of the response from the cue

under a decaying random walk through memory. The motivation for this procedure is that the similarity between two cue words might rely not just on the similarity of the responses they immediately activate, but also on the similarity of responses to *those* responses, and so forth.¹ To perform this transformation, first, each cue-response cell is subjected to the positive pointwise mutual information transformation, to decrease bias of the random walk toward highly connected responses (Newman, 2010). The activation of each response, in the limit of a decaying random walk, can then be obtained from the following transformation:

$$\mathbf{G}_{\text{rw}} = \sum_{r=0}^{\infty} (\alpha \mathbf{P})^r = (\mathbf{I} - \alpha \mathbf{P})^{-1}$$

where r is the step number in the random walk, P is the PPMI transformed association matrix, $\alpha < 1$ is a damping parameter, and I is the identity matrix (see De Deyne et al., 2019 for more explanation). The cosine similarity between cue vectors in such a matrix correlated well with semantic relatedness judgments from various norms and, importantly, similarity judgments in SimLex-999 with $r = .68$.

Despite the ostensibly similar purposes or capabilities of text-based and association-based representations, there have been relatively few direct comparisons of these two sources of representation. For example, De Deyne et al. (2015) obtained predictions of similarity and relatedness judgments among thematically related or co-hyponymic Dutch word pairs by taking cosine similarity between rows of a cue-response frequency matrix reweighted by positive

¹ We note that text-based models also often capture higher order co-occurrence relations. This is one reason why matrix factorization techniques like singular value decomposition are applied to word-context co-occurrence matrices, as in Latent Semantic Analysis (Landauer and Dumais, 1997; Jones et al., 2015; Schlechtweg, et al., 2019). Word2vec is also sensitive to second-order information (Schlechtweg, et al., 2019), likely because, as mentioned, it can be shown that word2vec implicitly performs matrix factorization (Levy & Goldberg, 2014).

pointwise mutual information (De Deyne et al., 2013). Pearson correlations between these predictions and actual judgments were substantially higher (advantages of $\sim r = .1$ to $r = .26$) than correlations between actual judgments and predictions obtained from cosine similarity between vectors from a model based on word-word co-occurrences in syntactic dependencies (e.g., *noun-verb* dependencies like “We need some more *coffee*”). Similarly, Vankrunkelsven et al. (2018) found that affective word properties (e.g., valence and arousal) of English and Dutch words were better predicted from k -Nearest Neighbors regression of such a PPMI-transformed cue-response matrix than they were from k -NN regression of a similar syntactic-dependency-based text model. One of the goals of the present work, therefore, is to examine whether representations from free association norms make better predictions of similarity judgments than do various representations based on lexical co-occurrence statistics in text, including several not considered by De Deyne et al. (2015) or Vankrunkelsven et al. (2018).

1.1.3 Representations for the present work

We consider six representations based on lexical co-occurrence statistics in text and three based on the De Deyne et al. (2019) free association norms. All representations have vectors for all the words in the eight sets of co-hyponyms we test. Our text-based representations are pre-trained by other authors and consequently vary simultaneously on training algorithm (and hyperparameters like window size in word2vec), corpus, and, if applicable, specialization method². Despite the difficulty this will sometimes create in understanding the independent effect of these factors on quality of similarity predictions, we use these representations for two

² One might object that this means we are not testing (representation, similarity function) pairs but rather (representation + training data, similarity function) or (algorithm + training data, similarity function) pairs. In our view, however, it is fair to say we are evaluating (representation, similarity function) pairs since representations are the output of algorithms operating on datasets (i.e., representation = algorithm(dataset)).

reasons. First, pre-trained representations are usually trained on very large corpora that are difficult for researchers without large computing resources (such as ourselves) to deal with. We expect that in the present setting where very fine-grained semantic relations must be respected, corpus size will be an important factor. Second, these pre-trained representations are often what cognitive scientists or natural language processing researchers actually use in their own applications (often because, again, it is difficult or impossible to train one's own vectors). We can therefore offer practical advice regarding the representations that researchers ought to use if they need estimates of word-word similarity. Further, we will still be able to make some key comparisons between representations trained on text. For one, we can evaluate the specialization techniques of Wieting et al. (2015) and Ponti et al. (2020), as we compare these to unspecialized embeddings trained on the same corpora (Gigawords plus Wikipedia, and Common Crawl, respectively), with the same training algorithm (GloVe). We can also evaluate the effect of corpus size by comparing, for example, GloVe trained on 6B words of GigaWords and Wikipedia, to GloVe trained on 840B words of Common Crawl. Researchers interested in comparing the training algorithms on the same corpora can examine the work of Levy et al. (2015), as well the initial papers introducing the representations (e.g. Pennington et al., 2014).

Of the three vector representations from association norms we test, the first representation is exactly the random walk vectors described above. Because these representations are of especially high dimensionality (12,216 dimensions), we also test these same vectors reduced to 300 dimensions via SVD. In particular, we take the entire 12,216 by 12,216 matrix composed of random walk vectors, apply truncated SVD, and retain the first 300 components. We do this both to determine whether the quality of these random walk vectors depends on the

high dimensionality, and because for some applications researchers may be interested in, the full 12,216 dimension representations may be computationally difficult to work with. Finally, we also use De Deyne et al.'s (2019) PPMI vectors reduced via SVD to 300 dimensions. Comparing these to the other two free association vectors allows us to quantify the value added by the random walk transformation. Further, because word2vec (SGNS) is equivalent to factorization of a PMI word-context matrix, comparing these vectors to word2vec will allow us to understand the different similarity-relevant information in natural language corpora and free association norms. Table 1 summarizes all nine representations we used.

Table 1. Vector representations evaluated in the present study.

Representation name	Training algorithm	Corpus	Corpus Size	Dimensionality	Specialization method	Source
word2vec 100B GoogleNews	Skip-gram with negative sampling (SGNS)	Google News	100B	300	None	Magnitude Library ^a
fastText 600B Common Crawl	FastText with Continuous Bag of Words (CBOW)	Common Crawl ^g	600B	300	None	fastText website ^b
GloVe 6B Giga+Wiki	GloVe	GigaWord ^h Corpus and Wikipedia	6B	300	None	GloVe website ^c
GloVe 6B Giga+Wiki Adversarial postspecialization	GloVe	GigaWord Corpus and Wikipedia	6B	300	Ponti et al. (2020)	First author website ^e
GloVe 840B Common Crawl	GloVe	Common Crawl	840B	300	None	Magnitude Library
GloVe 840B Common Crawl, Paragram	GloVe	Common Crawl	840B	300	Wieting et al. (2015)	First author website ^d
SWOW-RW	Random walk transformation of free association cue-response matrix	n/a	n/a	12,216	None	First author, personal comm ^f
SWOW-RW-SVD	SWOW_RW vectors reduced to 300 dimensions via SVD	n/a	n/a	300	None	Derived
SWOW-PPMI-SVD	SWOW PPMI vectors reduced to 300 dimensions via SVD	n/a	n/a	300	None	First author, personal comm ^f

^aThe magnitude library (Patel et al., 2018) can be found at <https://github.com/plasticityai/magnitude>.

^bThe fastText website can be found at <https://fasttext.cc/docs/en/english-vectors.html>.

^cThe GloVe website can be found at <https://nlp.stanford.edu/projects/glove/>.

^dParagram embeddings can be found at <https://www.cs.cmu.edu/~jwieting/>.

^eGloVe embeddings specialized with Ponti et al.'s (2020) method can be found at <https://github.com/cambridgeltl/adversarial-postspec>.

^fSWOW vectors were obtained by emailing the first author of De Deyne et al. (2019).

^gCommon Crawl is a corpus generated by crawling billions of web pages.

^hGigaWord is a comprehensive archive of newswire texts.

Note that in addition to the vector representations derived from text or free association, semantic representations are also sometimes obtained using feature norms, i.e. by explicitly asking participants to verbally generate features for objects (McRae et al., 2005; Buchanan et al.,

2019). Lexical representations from feature norms have been successfully applied to phenomena including at least semantic priming, lexical access, and similarity judgment (Cree et al., 1999; O'Connor et al., 2009; Peterson & Griffiths, 2017). Although these are possible representations, we will not evaluate these in the present work, as available feature norms do not sufficiently cover the items for which we wish to make predictions. For example, only 147 of the 195 words in our stimuli are in the Buchanan et al. (2019) norms, which are the largest feature norms we are aware of (with over 4000 entries).

1.2. Similarity functions

1.2.1 Unweighted functions

Cosine similarity is the similarity function most commonly applied to vector representations from text or association norms (Jones et al., 2015; Lenci, 2018; Jurafsky and Martin, in prep). It is defined as the dot product of two vectors $\mathbf{u}$ and $\mathbf{v}$ of length k , L2-normalized by the product of their lengths: $(\mathbf{u} \cdot \mathbf{v}) / (||\mathbf{u}|| * ||\mathbf{v}||)$. Cosine similarity is bounded between +1 and -1, with +1 corresponding to vectors that are in identical directions and -1 corresponding to vectors that are in opposite directions (vectors that are orthogonal have a cosine similarity of zero).

The unnormalized dot product itself, $\mathbf{u} \cdot \mathbf{v}$, also has some precedent as the similarity function at the heart of the common features or additive clustering models of similarity (Shepard & Arabie, 1979; Navarro & Lee, 2004). In such models, representations are essentially like those derived from feature norms, often represented as bit vectors where an element is 1 if an item possesses a feature and 0 otherwise. Although the dot product's history in psychological modeling largely lies in binary, featural representations, it can of course be applied to real-valued

vector representations like those derived from text or association norms. Also note that the dot product in the case of the common features or additive clustering models is usually applied using feature weights which reflect the importance of each feature in similarity computation --- we discuss the issue of weights later in this section.

Other similarity functions are also possible, and some even have deep theoretical precedent in psychological models of similarity. For example, in multidimensional scaling models of similarity (MDS; Shepard, 1962; Kruskal, 1964; Hout et al., 2013), the spatial model of similarity that in some ways anticipated the text- or association-based vector representations described above, *dissimilarity* is usually calculated according to Euclidean distance or city-block distance. Euclidean distance is defined as $\sqrt{\sum_k |u_k - v_k|^2}$; city-block distance is defined as $\sum_k |u_k - v_k|$. The generalization of these two functions, where the exponent of the elementwise difference and the rooting of the sum are determined by a parameter p , is known as the Minkowski distance. In most applications of multidimensional scaling, however, p is set to 1 (city-block distance) or 2 (Euclidean distance). Traditionally, $p = 1$ has been used for ‘separable’ stimuli dimensions, like size and shape, that can be attended to separately, while $p = 2$ has been used for ‘integral’ dimensions, like hue and saturation, that cannot be so easily separated (Garner, 1974). Finally, note that Euclidean distance and cosine similarity are related: squared Euclidean distance between L2-normalized vectors is proportional to cosine similarity.

Our last two similarity functions are Pearson correlation and Spearman correlation. While Pearson and Spearman correlation have somewhat less precedence as measures of similarity in psychological models of similarity or in applications with vector representations

(but see the COALS model of Rohde et al., 2006), they are closely related to the cosine similarity and dot product metrics discussed above.

We noted that cosine similarity has dominated in applications with vector representations, despite the *a priori* reasonableness of the alternative similarity functions discussed so far. It is true that there are some theoretical reasons for using cosine over alternatives. For one, although the details depend on the process generating the vectors, Euclidean distance can become uninformative in high-dimensional spaces (as typical word vectors derived from text or association norms are), as every vector becomes distant from every other vector³ (Aggarwal et al., 2001). In addition, for many or most text-based models, word frequency is correlated with vector length (distance from the origin) simply because a more frequent word has higher counts in its row of a word-context co-occurrence matrix. As it is usually simply assumed that word frequency is not considered during similarity judgment (Jurafsky & Martin, in prep), cosine is used because it normalizes for word vector length and hence ignores word frequency (cf. dot product, which is sensitive to vector length and hence frequency).

However, these predictions or assumptions, and the possible equivalence or even superiority of alternative similarity functions, are largely untested with empirical judgments of similarity. Pereira et al. (2016) conducted one of the few such systematic comparisons, finding that word2vec and GloVe performed better than older word representations (e.g., Latent Semantic Analysis; Landauer & Dumais, 1997), and that cosine and Pearson correlation performed better than Euclidean distance, at predicting similarity judgment from SimLex-999 [as

³ Another way of stating this problem is that, in high-dimensional spaces, the ratio of the Euclidean distances of the nearest and farthest neighbors to a given target is almost 1. That means that (dis)similarity is not tracked by Euclidean distance in high-dimensional spaces.

well as behavior in free association tasks (Nelson et al., 2004) and relatedness judgments (Bruni et al., 2012)]. Of course, as mentioned before, SimLex-999 is in some respects an undemanding benchmark of similarity, as it contains many cross-category word pairs that can be easily judged without nuanced comparisons between the items in the same category. In addition, Pereira et al. (2016) did not consider several of the similarity functions (Spearman correlation, (augmented) city block distance, dot product) or representations (text-based representations specialized for similarity as well as representations based on association norms) we have mentioned, nor did they consider weighted similarity functions.

Perhaps the closest test to what we have in mind comes from Bullinaria and Levy (2007), who compared various similarity functions' and text representations' performance on the synonym portion of the Test of English as a Foreign Language (TOEFL). Bullinaria and Levy found that cosine similarity consistently beat city block distance and Euclidean distance on this task (and other semantic tasks). However, this issue deserves further examination. For one, more advanced word vector representations have emerged in the 13 years since Bullinaria and Levy's work (Mikolov et al., 2013; Pennington et al., 2014; Wieting et al., 2015; Mikolov et al., 2018; Ponti et al., 2020), and these newer representations may favor different similarity functions. In addition, much like SimLex-999, the synonym portion of the TOEFL is not as challenging to model as the fine-grained, within-category similarity data many cognitive scientists are concerned with. Therefore, the tests we propose – modeling similarity judgments between hundreds of pairs of closely related co-hyponyms of specific categories – constitute stricter evaluations of various word vector representations and similarity functions.

1.2.2. Weighted functions

Of course, all of the similarity functions we have described so far treat all the dimensions of the representation equally. That is, the functions lack weights on the dimensions. Yet, there is good reason to believe that the absence of weights underestimates the utility of word vector representations for modeling human judgments of similarity, at least when considering similarity among pairs of cohyponyms or category coordinates. For one, variable weights on dimensions of (latent) item representations are key to several well-known cognitive models, from the weighted-additive decision rule in multi-attribute choice (Keeney et al., 1993), to the generalized context model of categorization (Nosofsky, 1984, 1986), which of course involves implicit judgments of similarity between novel and stored exemplars. But even more germane is work by Peterson et al. (2018). Peterson et al. collected Likert-scale similarity judgments for pairs of images from within six categories (animals, automobiles, fruits, furniture, vegetables, and a category of miscellaneous images), and then extracted real-valued vectors for these images using the hidden layers of deep neural networks (DNN's) trained to perform image classification. They found that the unweighted dot product between two images' vectors accounts for about 40% of the variance in human judgments of similarity. However, out-of-sample predictions of similarity could be much improved, accounting for 60% of the variance, by regressing human judgments onto a vector constructed from the elementwise product of two image vectors. In other words, by using human similarity judgments as supervision, Peterson et al. effectively learned a weighted dot product function (with weights potentially reflecting the dimensions or features that people preferentially attend to while evaluating similarity). Most striking, Peterson et al. showed that a weighted function could be fit to similarity judgments from many categories (e.g., animals, automobiles, fruits, and furniture) and make accurate predictions for a held-out

category (e.g., vegetables) above the predictions made by the unweighted dot product, suggesting that the learned weights were largely category-general. However, predictions from these category-general models were not quite as good as those obtained from training and testing a regression model on a single category, suggesting that the dimensions important for evaluating similarity in images from the categories Peterson et al. studied are partly, but not entirely, idiosyncratic to a category.

In contrast to the findings of (largely) category-general weights in Peterson et al. (2018), there is a large body of psychological literature suggesting that similarity judgment is highly context-dependent (Tversky, 1977; Barsalou, 1982; Medin et al., 1993; Goldstone et al., 1997; Close et al., 2009). For example, the well-known shape bias in noun learning in children can also be understood as the effect of linguistic context on similarity perception: when shown an object and then asked to “find another dax”, children are more likely to choose an object similar in shape, as opposed to size or texture, than they are if they’re asked to “find another one” (Landau et al., 1988). In Goldstone et al. (1997), a majority of participants demonstrated choice intransitivities in a 2AFC ‘which object is most similar to the probe?’ task, with a plurality adhering to a ‘diagnostic-dimension strategy’ that emphasized the single dimension that distinguished the two choices most clearly.

The context effect perhaps most relevant for the present work is the *extension effect*, according to which features influence similarity judgments more when they vary within an entire set of stimuli (Garner, 1962; Bransford et al., 1989; Pomerantz & Lockhead, 1991). For example, in one classic Tversky (1977) experiment, participants in one group rated the similarity of European countries, another group rated the similarity of South American countries, and a third

group rated similarity of two South American countries on some trials and two European countries on other trials. Similarity ratings from this last group were much higher than the ratings from the other two groups, presumably because features like “South American” or “European” became important when they were not held constant across the entire set of stimuli. Thus, to the extent that there are different dimensions of variation within different categories of words like professions and vehicles – the former varying on salary and required education and the latter varying on size, cost, and speed – dimension weights may vary by category. This phenomenon may be stronger with words than it is with images, where features like shape, color, and texture may be important cues to similarity regardless of category. In any case, it is not known if weights, category- or context-dependent or not, can even be effectively learned to improve out-of-sample predictions of similarity from word vector representations. Determining this is therefore another aim of the present work. For every unweighted function we have mentioned above, we can consider a weighted counterpart – see Table 2 for all unweighted and weighted functions we consider.

We consider one additional weighted similarity function without an unweighted counterpart, derived from recent work in relation learning with word vector representations, which may be of use in addressing the potential category- or context-dependence of similarity judgment. In the relation learning setting, the problem is to determine if a word pair possesses a particular lexical relation like ISA (cat-ISA-animal), CAUSES (smoking-CAUSES-cancer), or PARTOF (wheel-PARTOF-car). To solve this problem, Lu et al. (2019) extended their earlier BART model (Bayesian Analogy with Relational Transformations; Lu et al., 2012; also see Chen et al., 2017). BART has multiple stages and components, but its core idea is to use logistic

regression to learn a mapping to relation judgments from the concatenation of (a) the raw difference vector between two words, and (b) the *sorted* difference vector. Lu et al. (2019) explain that the function of sorting is to “highlight semantic features that tend to be aligned with respect to functionally relevant differences between the two words in a pair” (pg 4177). Indeed, this procedure appears to work well, as out-of-sample predictions achieve correlations with human typicality judgments for relations that approach inter-rater reliability. Thus, the final weighted similarity function we consider is the core component of BART outlined here (see Table 2), which we refer to as *augmented* city-block distance.

Having introduced all our unweighted and weighted similarity functions, we now point out that our similarity functions generally fall into two groups: cosine, dot product, and Spearman and Pearson correlation in one group, and Euclidean, city-block, and augmented city-block distance in another. As already mentioned in the previous section, the former group, is closely related to classic common features psychological models (Shepard & Arable, 1979; Navarro & Lee, 2004). In these models, similarity increases when two items share a feature, or similar values on a dimension. The latter group, in contrast, is consistent with distinctive features models (Navarro & Lee, 2004). In these models, similarity *decreases* when one item has a feature the other doesn’t, or when two items have different values on a dimension. This distinction is possibly relevant for the present work, as some past work has shown that common features matter more than distinctive features for similarity judgments in linguistic stimuli (Gati & Tversky, 1984; Ritov et al., 1990; Goldstone & Son, 2004). We therefore might expect our common features models to dramatically outperform the distinctive feature models. We also point that, as with representations, different similarity functions constitute different

computational level theories (Marr, 1982) of *what* is being computed during similarity judgment, as opposed to *how* that computation is achieved. If the best computational level theory of human similarity judgment could be determined, this could constrain future work about how similarity is computed at the algorithmic or implementational levels.

Table 2. Unweighted and weighted similarity functions evaluated in the present study.

Function name	Unweighted ^a	Weighted ^a
Dot product	$\mathbf{u} \cdot \mathbf{v} = \sum_k \mathbf{u}_k * \mathbf{v}_k$	$\sum_k \mathbf{w}_k * \mathbf{u}_k * \mathbf{v}_k$
Cosine	$(\sum_k \mathbf{u}_k * \mathbf{v}_k) / (\mathbf{u} * \mathbf{v})$ where $ \mathbf{u} = \sqrt{(\sum_k \mathbf{u}_k ^2)}$	$\sum_k \mathbf{w}_k * \mathbf{u}_k * \mathbf{v}_k$ after vectors $\mathbf{u}$ and $\mathbf{v}$ have been L2-normalized, so that each word vector's length is 1
City block distance	$\sum_k \mathbf{u}_k - \mathbf{v}_k $	$\sum_k \mathbf{w}_k * \mathbf{u}_k - \mathbf{v}_k $
Euclidean distance ^b	$\sqrt{(\sum_k \mathbf{u}_k - \mathbf{v}_k ^2)}$	$\sum_k \mathbf{w}_k * \mathbf{u}_k - \mathbf{v}_k ^2$
Pearson correlation	$\sum_k \mathbf{u}_k * \mathbf{v}_k$, after vectors $\mathbf{u}$ and $\mathbf{v}$ have been z-scored	$\sum_k \mathbf{w}_k * \mathbf{u}_k * \mathbf{v}_k$, after vectors $\mathbf{u}$ and $\mathbf{v}$ have been z-scored
Spearman correlation	$\sum_k \mathbf{u}_k * \mathbf{v}_k$, after vectors $\mathbf{u}$ and $\mathbf{v}$ have been converted to ranks and then z-scored	$\sum_k \mathbf{w}_k * \mathbf{u}_k * \mathbf{v}_k$, after vectors $\mathbf{u}$ and $\mathbf{v}$ have been converted to ranks and then z-scored
Augmented city block distance ^c	n/a	$\sum_k \mathbf{w}_k * \mathbf{u}_k - \mathbf{v}_k \oplus \sum_k \mathbf{b}_k * \mathbf{u}_k - \mathbf{v}_k _{\text{sorted}}$

^a All functions compute similarity between two vectors $\mathbf{u}$ and $\mathbf{v}$ of dimensionality k . Weighted functions introduce a weight vector $\mathbf{w}$ over all k dimensions.

^b In the weighted case of Euclidean distance, we do not take the square root of the sum over all k dimensions, as would ordinarily be done when computing euclidean distance. We do this so that we can use off-the-shelf linear regression packages to find the weights. This choice of course does not change the ordering of the resulting similarity predictions, which is the only thing our metric of performance, Spearman correlation, is sensitive to.

^c The augmented city block function also has a weight vector $\mathbf{b}$ over the sorted difference vector; the index k for these weights and the sorted difference vector $|\mathbf{u}_k - \mathbf{v}_k|_{\text{sorted}}$ can be understood as indexing the dimension with the k th greatest difference between vectors $\mathbf{u}$ and $\mathbf{v}$.

1.3 Present study

In sum, the goal of the present study was to investigate the ability of various general-purpose lexical representations derived from text and free association norms to model human judgments of similarity among co-hyponyms of various categories, using various similarity metrics, which are all related to classic psychological models of similarity like common features/additive clustering, distinctive features, and multidimensional scaling. We first assessed the ability of these representations and the *unweighted* versions of various similarity functions to predict human judgments of co-hyponym similarity. We then used the similarity judgments to

supervise learning of weights on the dimensions, for each candidate representation and similarity function. We first learn weights on a per-category basis, and then we attempt to learn weights that generalize across several categories or contexts at once.

2. Methods

2.1 Words

We collected similarity data for eight categories: furniture, clothing, birds, vegetables, sports, vehicles, fruit, and professions. Each category contained 20-30 words referring to category members (see Appendix A). These categories and items were loosely inspired by the Leuven Concept Database (De Deyne et al., 2008), which contains similarities for such categories in Dutch.

2.2 Participants

We recruited 365 participants (mean age = 33 years, 55% female) through Prolific Academic. We limited our data collection to participants who were from the U.S. and had an approval rate above 80%. Participants were only allowed to participate once, and were paid approximately \$10 per hour.

2.3 Design and Procedure

We used a between-subjects design whereby participants were randomly assigned to one of the eight categories – furniture ($N = 33$), clothing ($N = 61$), birds ($N = 54$), vegetables ($N = 30$), sports ($N = 61$), vehicles ($N = 28$), fruit ($N = 31$), and professions ($N = 67$). Twice as many participants were required for birds, clothing, professions, and sports because each participant randomly assigned to those categories only completed pairwise judgments for half the items (a choice we explain shortly).

Participants were told that we were interested in how people judge similarity of word meaning, on a Likert scale from 1 (not at all similar) to 7 (extremely similar). Participants were given three examples of words pairs (from categories not tested here) for each of three different levels of similarity: high, medium, and low. We then (1) showed participants a page with the name of the category they'd be rating along with each of the 20-30 words, (2) indicated how many pairs in total they would be rating, and (3) told participants they would be rating similarity in blocks of at most 50 pairs on a single page. Our intention behind showing participants the name of the category as well as all the items was to help participants calibrate their similarity scale and think about the dimensions or features that could distinguish these items as a whole. Similarly, having 50 pairs on a page would allow participants to calibrate their response to one pair based on their responses to other pairs. Our hope was that these design choices would mitigate previous concerns with pairwise Likert scale ratings of similarity that participants' strategies for rating change partway through the survey (Johnson et al., 1990; Hout et al., 2013). After this preview of the survey, participants proceeded to the first block.

Pairs of words were generated to cover all possible pairs and such that every word occurred equally often in the first position and the second position. Crucially, we did not collect each pair of words twice, reversing the order of the words. Pairs were then randomly assigned once (for all participants) to a block, and order of pairs within block and order of blocks was randomized differently for each participant. For the categories birds, clothing, professions, and sports, participants were randomly assigned to only half of all the blocks, as completing all blocks would have taken a prohibitive amount of participant time and increased chances of burnout (see category member counts in Appendix A). After completing all blocks, participants

indicated their age and gender. All behavioral data we collected can be found in an OSF repository for a related project:

https://osf.io/d7fm2/?view_only=c5ba5d34a5e34ff3970a652c07aadc5c.

3. Results

	(A) Unweighted								(B) Per-category weights							
w2v 100B GN	39	35	39	38	39	39	40	38	66	66	66	67	63	63	65	65
ft 600B CC	53	45	53	51	50	50	53	51	72	70	72	71	71	71	73	71
GI 840B CC	51	48	51	50	49	49	49	50	72	72	73	73	72	71	71	72
GI 840B CC, Paragram	45	45	45	45	32	32	33	40	61	61	61	61	56	56	58	59
GI 6B Giga+Wiki	40	38	40	40	37	37	37	39	65	64	64	65	64	63	64	64
GI 6B Giga+Wiki, AP	39	39	39	40	39	39	39	39	60	60	61	60	61	62	63	61
SWOW-RW	57	43	57	45	52	52	52	51	64	77	77	75	77	77	77	75
SWOW-RW-SVD	55	42	55	61	57	57	51	54	68	74	73	74	74	73	73	73
SWOW-PPMI-SVD	61	53	61	60	55	55	53	57	75	74	75	75	71	71	72	73
Average	49	43	49	48	46	46	45	47	67	69	69	69	68	68	69	68
	(C) Hold one category out								(D) Mix of categories							
w2v 100B GN	37	37	37	38	34	34	34	36	51	50	50	50	45	46	46	48
ft 600B CC	52	46	53	51	53	53	54	52	61	57	61	59	62	61	62	60
GI 840B CC	48	49	48	48	46	46	47	47	60	61	61	58	57	55	57	59
GI 840B CC, Paragram	43	42	43	42	29	29	29	37	49	48	49	47	39	39	40	44
GI 6B Giga+Wiki	37	36	37	36	32	32	32	35	49	48	48	47	42	40	43	45
GI 6B Giga+Wiki, AP	35	35	35	35	36	36	36	35	43	43	43	43	45	44	46	44
SWOW-RW	28	41	48	53	50	49	47	45	57	68	70	68	70	71	73	68
SWOW-RW-SVD	57	58	59	58	55	55	55	57	63	65	67	65	63	63	64	64
SWOW-PPMI-SVD	59	57	59	59	53	53	52	56	65	64	65	65	60	61	61	63
Average	44	45	46	47	43	43	43	44	55	56	57	56	54	53	55	55
	Cosine similarity	Dot product	Pearson correlation	Spearman correlation	Augmented city block distance	City block distance	Euclidean distance	Average	Cosine similarity	Dot product	Pearson correlation	Spearman correlation	Augmented city block distance	City block distance	Euclidean distance	Average

Fig. 1. Performance of every (representation, similarity function) pair in the four evaluation settings. Cells indicate the Spearman correlation between mean human judgments and similarity function predictions, averaged over categories and multiplied by 100 (for readability). Note that, in unweighted cases (subfigure A), augmented city block distance reduces to city block distance. Also, in unweighted cases (subfigure A), we have flipped the sign of the Spearman correlations for the three distance functions to make them more comparable to the other four functions.

3.1 Unweighted similarity functions

In all analyses, we first average over all participants' similarity judgments to yield a single mean rating of similarity for a given word-word pair.⁴ Fig. 1A summarizes performance (averaged over all eight categories⁵) for each representation and similarity function. Here the columns correspond to similarity functions, the rows correspond to the underlying representation, and the cells indicate the Spearman correlation between averaged human judgments of similarity and the predictions of that representation and similarity function combination. Note that some functions, like dot product, Pearson and Spearman correlation, and cosine, are similarity functions, whereas Euclidean and city-block distances are *dissimilarity* functions (a distinction that, again, tracks the common features vs distinctive features classes of similarity models, Navarro & Lee, 2004).⁶ We have flipped the sign on the correlations for these latter two, so that they can be more easily compared to dot product, Pearson and Spearman correlation, and cosine.

The bottom row of Fig. 1A suggests that the best performance is obtained from text unspecialized vectors trained on the largest corpora, fastText-600B-CommonCrawl (mean $r = .51$) and GloVe-840B-CommonCrawl (mean $r = .50$), as well as vectors from free association norms (SWOW-RW mean $r = .51$, SWOW-RW-SVD mean $r = .54$, SWOW-PPMI-SVD mean $r = .57$), with only small differences among these. Text-based vector representations specialized to respect similarity relations appear to provide no or minimal advantage over their unspecialized

⁴ See the supplemental materials for an analysis of individual differences that justifies averaging. Modeling individual differences is left for future work

⁵ The supplemental materials contains a figure analogous to Figure 1, but without averaging over categories.

⁶ For consistency with other subplots in Fig. 1, we have also included the augmented city block distance, but note that this function reduces to the city block distance in the unweighted setting.

counterparts, as can be seen by comparing GloVe-840B-CC (mean $r = .50$) to GloVe-840B-CC-Paragram (mean $r = .40$) and GloVe-6B-Giga+Wiki (mean $r = .39$) to GloVe-6B-Giga+Wiki-Adv-Postspec (mean $r = .39$). In fact, the specialization procedure of the GloVe-840B-CC-Paragram vectors may have in fact hurt performance: the unspecialized vectors have approximately a $r = .1$ advantage in performance over the specialized vectors. Fig. 1A also suggests only small differences between similarity functions, with perhaps small advantages for cosine and Pearson correlation (mean r 's = $.49$) over Spearman correlation (mean $r = .48$), city block and Euclidean distance (mean r 's = $.45$) and dot product (mean $r = .43$).

Vehicles								Vegetables							
w2v 100B GN	70	62	70	71	71	71	69	30	31	30	27	29	29	27	
fT 600B CC	58	46	59	54	53	53	60	57	60	57	58	43	43	46	
GI 840B CC	55	56	55	56	50	50	49	39	27	39	39	41	41	42	
GI 840B CC, Paragram	67	70	67	69	38	38	38	38	32	38	36	41	41	42	
GI 6B Giga+Wiki	63	62	63	67	57	57	59	30	24	30	33	36	36	36	
GI 6B Giga+Wiki, AP	53	53	53	54	55	55	53	17	17	17	18	17	17	17	
SWOW-RW	71	67	72	60	70	70	68	40	5	39	32	39	39	43	
SWOW-RW-SVD	70	66	70	84	74	74	67	37	4	37	49	45	45	37	
SWOW-PPMI-SVD	74	62	74	80	75	75	71	47	30	47	52	49	49	47	
Sports								Clothing							
w2v 100B GN	34	28	34	34	44	44	45	44	39	43	42	44	44	45	
fT 600B CC	41	33	42	38	48	48	53	64	56	64	64	63	63	64	
GI 840B CC	43	40	43	39	49	49	49	61	55	62	62	63	63	64	
GI 840B CC, Paragram	38	39	38	38	32	32	32	53	53	53	54	45	45	46	
GI 6B Giga+Wiki	32	29	32	27	34	34	36	37	36	37	39	38	38	37	
GI 6B Giga+Wiki, AP	35	35	35	37	36	36	35	45	45	45	45	42	42	45	
SWOW-RW	52	55	56	51	46	46	43	58	43	57	38	52	52	57	
SWOW-RW-SVD	50	55	50	53	48	48	42	58	42	58	73	66	66	56	
SWOW-PPMI-SVD	57	55	57	53	45	45	45	67	58	67	73	66	66	61	
Birds								Fruit							
w2v 100B GN	21	19	21	20	16	16	16	28	29	29	29	22	22	22	
fT 600B CC	38	35	38	36	31	31	34	49	45	49	46	44	44	46	
GI 840B CC	41	39	41	43	38	38	38	44	43	44	48	42	42	41	
GI 840B CC, Paragram	22	21	22	21	17	17	17	34	33	34	37	35	35	35	
GI 6B Giga+Wiki	20	18	20	20	23	23	21	37	36	37	38	35	35	37	
GI 6B Giga+Wiki, AP	12	12	12	12	11	11	12	39	39	39	40	38	38	39	
SWOW-RW	45	28	46	31	38	38	42	54	17	53	32	50	50	53	
SWOW-RW-SVD	43	25	43	51	45	45	40	54	15	54	63	58	58	51	
SWOW-PPMI-SVD	42	32	42	41	39	39	37	60	45	61	67	61	61	58	
Professions								Furniture							
w2v 100B GN	42	38	42	41	42	42	42	43	37	43	43	46	46	49	
fT 600B CC	51	36	51	47	54	54	53	63	51	63	63	65	65	67	
GI 840B CC	63	65	63	50	51	51	53	63	63	63	65	62	62	58	
GI 840B CC, Paragram	49	54	48	46	10	10	11	61	61	60	58	39	39	43	
GI 6B Giga+Wiki	48	52	48	47	26	26	24	50	49	50	53	50	50	46	
GI 6B Giga+Wiki, AP	50	50	50	49	48	48	50	65	65	65	65	65	65	65	
SWOW-RW	62	66	66	60	52	52	47	70	61	70	55	68	68	67	
SWOW-RW-SVD	61	66	61	43	48	48	46	70	61	70	73	74	74	67	
SWOW-PPMI-SVD	64	69	64	45	38	38	37	75	73	75	67	67	67	67	
Cosine similarity								Cosine similarity							
Dot product								Dot product							
Pearson correlation								Pearson correlation							
Spearman correlation								Spearman correlation							
Augmented city block distance								Augmented city block distance							
City block distance								City block distance							
Euclidean distance								Euclidean distance							

Fig. 2. Performance of every (representation, similarity function) pair in the unweighted setting. Cells indicate the Spearman correlation between mean human judgments and similarity function predictions, multiplied by 100 (for readability). Note that augmented city block distance reduces to city block distance. Finally, we have flipped the sign of the Spearman correlations for the three distance functions to make them more comparable to the other four functions.

Considering average performance of a vector type over various similarity functions (or vice versa), however, can hide the possibility that particular similarity functions are well- or ill-suited for particular vector representations. Fig. 1A also shows that, in fact, for some representations, like GloVe-6B-Giga+Wiki-Adv-Postspec, similarity function matters little. For other representations, like SWOW-RW, similarity function matters a fair amount. For the latter representation, Pearson correlation has a $r = .14$ advantage over dot product. It can be seen that generally, vectors from association norms (SWOW-RW, SWOW-RW-SVD, and SWOW-PPMI-SVD) may have a small advantage (up to $r = .08$) over the best (text-based vector, similarity function) pairs, when similarity functions apparently unsuitable for a particular vector representation are discounted (i.e., Euclidean distance for SWOW-RW-SVD or dot product for SWOW-RW). Overall, SWOW-RW-SVD vectors using Spearman correlation as a similarity function, and SWOW-PPMI-SVD with cosine or Pearson correlation as similarity functions, appear to have the best overall performance ($r = .61$), beating SWOW-RW with Pearson correlation as the similarity function by almost .04 correlation points, and beating the best (text-based representation, similarity function) pair (fastText-600B-CC, Pearson correlation) by more than .08 correlation points.

Fig. 3. Split-half reliability for each of the eight categories we tested (green bars). Reliabilities result from computing Spearman correlation for many random splits. In blue is the best (representation, unweighted similarity function pair) from section 3.1. In orange is performance of the best (representation, weighted similarity function) pair, trained on a per-category basis as in section 3.2. Error bars indicate 95% credible intervals for predicted vs actual correlations, and the middle 95% of the distribution of split-half correlations.

While the above results give some insight into the representations and/or similarity functions that generally perform the best across all categories, it is also worthwhile to examine how performance depends on category. Fig. 2 contains heatmaps of Spearman correlations between mean human judgments and representation-function predictions for all categories, representations, and functions. A few notable patterns are visible. First, performance seems to vary quite a bit by category, with similarity among birds being predicted on average at $r = .29$ and vehicles at $r = .63$. Second, this ranking of categories by predictability seems to be fairly consistent across representations and similarity functions; for example, the heatmap for vehicles is generally ‘brighter’ than that for birds, reflecting the former category’s higher predictability. Third, there are a handful of cells in the heatmaps of Fig. 2 indicating extremely poor

performance for categories that are otherwise predictable and for representations and similarity functions that are otherwise predictive. For example, dot product with SWOW-RW or SWOW-RW-SVD makes especially poor predictions of similarity among fruits and vegetables, even though (a) this category is reasonably predictable for all other representation-metric pairs, and (b) dot product with SWOW-RW-SVD works fairly well for most other categories.

Finally, to better understand both between-category differences, as well as establish rough upper bounds on each category's predictability from vector representations, we computed split-half reliability for each category. To do this, we randomly split every category's set of participants in half, averaged over participants' similarity ratings within each half, calculated the Spearman correlation between each half's similarities, and applied the Spearman-Brown correction (Brown, 1910; Spearman, 1910). We did this for 100 random splits, yielding a distribution of correlations for each category. Fig. 3 contains barplots of these distributions for each category (along with the best unweighted and weighted (representation, similarity function) pairs). As can be seen, split-half reliability was quite high, with each category's average corrected correlation above $r = .8$, with an average across categories of $r = .89$.⁷ That these reliabilities are much higher than our best (representation, unweighted similarity function) pair ($r = .61$) suggests that there is substantial predictable variation in similarity judgments that we have not modeled. In addition, variation in category-level reliability seems to at least partially explain variation in predictability of similarity judgments from vector representations; category-level split-half reliability may be modestly correlated with average category performance according to a Bayesian correlation analysis with Morey and Rouder's (2018) R package *BayesFactor*

⁷ We do not focus on category differences in reliability. Such differences are likely to have many sources – subjects may be torn between folk and scientific taxonomies for birds, for example. As we said in a previous footnote, we leave further study of individual differences to future work.

[posterior median $r = .34$, 95% credible interval $(-.19, .75)$]. In the supplemental materials, we also show that word frequency and familiarity may have positive relationships with performance (for text-based vector representations and association vectors, respectively), but that these effects are highly uncertain given the small number of categories.

3.1.1 Interim discussion

A few key results are worth reiterating and briefly discussing. First, and perhaps least surprising, representations trained on larger corpora predicted similarity judgments better. This is unsurprising to the extent that larger corpora provide more contexts of occurrence for our words, and hence more information with which word vectors are specified. This may be especially important with the present similarity judgments, which arguably require more fine-grained representations than the judgments in past datasets like SimLex-999. It is also unsurprising that, when comparing representations trained on very large, roughly equally sized corpora (600 vs 840 billion words), training algorithm – fastText of GloVe – matters little. This is consistent with work showing that when parameters (e.g. context window size) of count models like GloVe are carefully tuned, there is no significant difference between these and prediction models like fastText (Levy et al., 2015), although of course we made no such attempt to tune model parameters, instead using off-the-shelf representations.

More surprising is that despite its somewhat greater popularity, cosine presented only small advantages over other similarity functions, advantages which vary somewhat by representation. Pearson correlation, which has comparatively little precedence in applications involving word vectors or spatial models of similarity more generally, performed just as well as cosine (likely because of their close functional relationship). It is also perhaps somewhat

surprising that Euclidean distance performed nearly as well as cosine, given the mathematical concerns, mentioned in the introduction, with Euclidean distance in high-dimensional spaces (Aggarwal et al., 2001).

Likewise, despite the success of similarity-specialization procedures (Glavaš et al., 2019) for improving performance on proper similarity judgments (as opposed to relatedness judgments) in benchmarks like SimLex-999, these techniques provide no advantage for the present similarity judgments among co-hyponyms (cf. the pairs of words drawn from the entire lexicon in, e.g., SimLex-999). The specialization method involved with the popular Paragram-SL999 vectors (Wieting et al., 2015) may have actually degraded performance in the present setting. As we observe similar results in the following sections, we save interpretation of the relative performance of specialized and unspecialized text-based vectors for the General Discussion.

Finally, considering the best similarity functions for a given representation, vectors based on free association norms presented small advantages over vectors based on text, an effect consistent with previous work comparing these two sources of representation when modeling semantic judgments (De Deyne et al., 2015; Vankrunkelsven et al., 2018). Even more striking was that Spearman correlation, when combined with SWOW-RW-SVD vectors, tied with cosine or Pearson correlation with SWOW-PPMI-SVD for the best overall performance ($r = .61$), .04 points better than SWOW-RW with Pearson correlation, and .08 points better than the best (text-based vector, function) pair. It is of course fairly well-known that dimensionality reduction with SVD can improve the quality of *text-based* vector representations. For example, Landauer & Dumais (1997) found that performance on the synonym portion of the TOEFL with latent

semantic analysis peaked around 300 dimensions retained via SVD, decreasing with fewer or more dimensions. It is perhaps interesting to see similar results with association-based vectors, *for certain similarity functions*, i.e., city block distance and Spearman correlation. Lastly, the fact that SVD-reduced association vectors performed *better* than unreduced SWOW-RW vectors suggests that this potential advantage of association norm vectors over text vectors is not due to the higher dimensionality of the association norm vectors, at least when using unweighted similarity functions. Rather, it seems there is some information in vectors from word association norms that is not present in the word vectors. We return to this issue in the General Discussion.

We also observed substantial variation between categories in terms of the predictability of their similarity judgments from vector representations, with more than .3 correlation points difference between the easiest category (vehicles) and the most challenging category (birds). This variation was somewhat explainable by variation in category-level split-half reliability; categories where participants agreed on the similarity judgments were easier for vector representations to predict. Yet this explained less than half of the variation between categories, and there was weak evidence for effects of frequency and familiarity as well. Better understanding the factors explaining category differences may therefore be an avenue for future work. It is also unclear, for now, why vectors from association norms combined with dot product make such extremely poor predictions of fruit and vegetables, even though (a) this category is reasonably predictable for all other representation-metric pairs, and (b) dot product with vectors from association norms works fairly well for most other categories.

So far, our work is somewhat similar to other work, e.g., Pereira et al. (2016) or Bullinaria and Levy (2007), who also tested various combinations of vector representations and

similarity functions on their ability to predict lexical-semantic behaviors, including free association, analogy, and similarity and relatedness judgments. We have already extended this work in two ways: (1) we used more recent text-based vector representations specialized for similarity (Wieting et al., 2015; Ponti et al., 2020) as well as random walk vectors based on free association norms (De Deyne et al., 2019), (2) we considered similarity functions not considered previously (e.g., Spearman correlation) (3) we evaluated representation-function pairs on a new, more challenging set of similarity judgments between co-hyponyms of various categories. What previous work has also not done, however, is examine (1) the possibility that weighted functions on text-based or association norm-based vectors, learned through supervision from human judgments of similarity, might make better predictions of similarity among words, and (2) how the weights of these functions might depend on context, category, or domain. We explore these possibilities in the next two sections.

3.2 Category-specific weighted similarity functions

Our general approach in this section is to learn weighted similarity functions by linearly regressing the similarity judgment between a pair of words on a vector that is the result of some elementwise operation on the vectors $\mathbf{u}$ and $\mathbf{v}$ for these two words (again, see Table 2 for all weighted similarity functions). As is standard in fitting weighted distance or similarity functions in previous work (see Navarro & Lee, 2004, when fitting, e.g., the common features model, or Nosofsky, 2014 when fitting the generalized context model of categorization, which embeds a weighted Minkowski distance), we constrain weights to be non-negative, so that the weights admit a straightforward interpretation of the ‘salience’ or ‘importance’ of each

dimension.⁸ Consequently, because Euclidean distance, city block distance, and augmented city block distance are *dissimilarity* and not similarity functions, we flip the sign of every element of the elementwise comparison vectors constructed as part of these functions. For example, when fitting the weighted city block distance, we derive the elementwise absolute value of the difference between two vectors $\mathbf{u}$ and $\mathbf{v}$, denoted $d_k = |u_k - v_k|$. This vector d_k necessarily contains all non-negative elements; we simply negate each element to convert d_k into a vector with all non-positive elements).⁹ In addition, because the vector representations are high-dimensional (≥ 300) relative to the number of observations we have for each category (190 for furniture and vegetables to 435 for birds), we employ L2-regularization to avoid overfitting. To simultaneously perform model selection (i.e., find the optimal regularization strength parameter α^{10}) *and* model evaluation, we employed nested cross-validation (Cawley & Talbot, 2010), with 3 inner folds (where we select α) and 3 outer folds (where we evaluate model predictions). Optimization of α was determined with Spearman correlation between predicted and actual judgments of similarity. All models were fit with the ElasticNet function in the scikit-learn machine learning package for Python (Pedregosa et al., 2011). To perform a pure L2-regularization, L1_ratio was set to 0. All other hyperparameters left to their defaults.

⁸ In other words, if we allowed and obtained negative weights, that would mean that two words were predicted to be *less* similar if they were *more* alike on a dimension with a negative weight.

⁹ We also fit an intercept for all weighted functions. The implication of this for, e.g., the weighted distance functions, is that one can think of the elementwise differences as subtracting from an initial positive and fixed amount of similarity. See also Navarro and Lee (2004) for a similar interpretation of the intercept in the context of the common features and distinctive features models of similarity.

¹⁰ We try every α that is an integer power of 10, from 10^{-6} to 10^6 .

Fig. 4. Performance of every (representation, similarity function) pair in the within-category weighted setting. Cells indicate the Spearman correlation between mean human judgments and similarity function predictions, multiplied by 100 (for readability).

Fig. 1B contains heatmaps summarizing performance averaged over the three outer folds, and over all eight categories of words. First, it is apparent that, for all vector types and similarity functions, performance is much improved relative to their unweighted counterparts (compare Figs. 1A and 1B). Excluding augmented city block distance which contained no analog in the unweighted setting (or rather, without weights, it is equal to twice the city block distance), we find an average improvement of $r = .22$. Second, we see largely the same ordering of vector types in terms of their performance: (a) text-based vectors from larger corpora outperform text-based vectors from smaller corpora, (b) similarity-specialized text-based vectors perform no better – and now, *consistently* worse across similarity functions – than their unspecialized counterparts, and (c) vectors from association norms outperform vectors from text. Second, with the introduction of weights, differences between similarity functions have diminished: all similarity functions achieve performance, averaged over categories and vector types, between $r = .67$ and $r = .69$. This invariance to similarity function holds within vector types, as well – the average difference between the best and worst similarity function for a representation decreased from .08 in the unweighted case to .04 in the weighted case. It is interesting to note, however, that the strongest exception to this pattern comes from the random walk vectors from free association norms, which somewhat disprefer the weighted cosine function, despite unweighted cosine’s popularity. Finally, the best (representation, weighted function) pair is the unreduced SWOW-RW vectors with a city block distance function, which now achieve a cross-validation

Fig. 5. Scatterplots of mean similarity judgments (x-axes) against predicted similarity judgments (y-axes) for the best (representation, weighted similarity function) pair: SWOW-RW with city block distance. Weights are learned separately for each category. Models are trained on 2/3 of the data for each category, and generate out-of-sample predictions for the remaining 1/3. Inset are Spearman correlations between predicted and actual judgments. The dotted line $y=x$ represents perfect predictions.

score of $r = .774$ (approximately .16 correlation points better than the best unweighted (representation, similarity function) pair), although we again emphasize that this is hardly better than other similarity functions operating SWOW-RW, which also achieve performance around $r = .77$. See Fig. 5 for scatterplots of predicted and actual judgments of similarity for each of the eight categories, for SWOW-RW vectors with a weighted city block distance function. Moreover, SWOW-RW vectors with a weighted city block distance function is still about .05 correlation points better than the best text-based (representation, similarity function) pair, fastText-600B-CC with squared difference vectors. Of course, this performance ($r = .774$) is still below split-half reliability, which, again, is above $r = .8$ for all eight categories, with an average of $r = .89$.

As with unweighted similarity functions, we also examined variation in performance by category. Fig. 4 contains heatmaps of correlations between predicted and actual judgments for

all categories, representations, and similarity functions. A few patterns are worth attention. First, the same categories that were easier to predict with unweighted models are easier to predict with weighted models (e.g. vehicles and furniture), although the difference between the easiest category (vehicles, $r = .80$) and the most difficult category (birds, $r = .57$) has shrunk somewhat (difference in r of .23, versus .33 in the unweighted case). Second, category-level split-half reliability is now an even stronger predictor of performance [posterior median $r = .59$, 95% credible interval (.13, .87)]. Finally, in contrast to Fig. 2 displaying results with unweighted functions, there are fewer anomalous cells like dot product and association norm vectors in the unweighted case.

Because we found very similar performance across all weighted similarity functions, we suspected that different functions were making very similar predictions (the alternative being that different functions made different predictions, but happened to capture very similarly sized, complementary proportions of variance). Thus, we checked the Spearman correlations between the predictions made by each of the weighted similarity functions. Fig. 6 displays these correlations for every vector type, averaged across categories. It is clear that there is generally a very high degree of correlation between different functions' predictions, with nearly perfect correlations within the common features (weighted dot product, cosine, Spearman and Pearson correlation) and within the distinctive features [(augmented) city block and Euclidean distance] groups introduced in the introduction. A model recovery analysis, explained in the supplemental materials, led to similar conclusions, suggesting that there are substantial identifiability issues among the weighted similarity functions we consider, at least for the stimuli and vector representations we have explored.

		w2v 100B GN							ft 600B CC							GI 840B CC						
Function2	Cosine similarity	100	97	99	97	79	81	81	100	98	100	96	91	91	93	100	98	100	96	89	89	89
	Dot product	97	100	96	94	67	70	69	98	100	98	96	88	88	92	98	100	97	94	83	84	84
	Pearson correlation	99	96	100	97	78	80	81	100	98	100	96	91	91	93	100	97	100	96	89	89	89
	Spearman correlation	97	94	97	100	79	81	80	96	96	96	100	90	90	92	96	94	96	100	88	88	88
	Augmented city block distance	79	67	78	79	100	99	97	91	88	91	90	100	99	95	89	83	89	88	100	99	97
	City block distance	81	70	80	81	99	100	97	91	88	91	90	99	100	96	89	84	89	88	99	100	97
	Euclidean distance	81	69	81	80	97	97	100	93	92	93	92	95	96	100	89	84	89	88	97	97	100
		GI 840B CC, Paragram							GI 6B Giga+Wiki							GI 6B Giga+Wiki, AP						
Function2	Cosine similarity	100	97	99	96	85	84	88	100	98	100	97	88	87	91	100	100	98	96	88	87	90
	Dot product	97	100	98	95	83	83	87	98	100	98	95	89	88	91	100	100	98	96	88	87	90
	Pearson correlation	99	98	100	95	85	84	88	100	98	100	97	88	87	91	98	98	100	97	87	86	90
	Spearman correlation	96	95	95	100	87	86	87	97	95	97	100	89	88	90	96	96	97	100	89	89	90
	Augmented city block distance	85	83	85	87	100	100	94	88	89	88	89	100	99	96	88	88	87	89	100	99	96
	City block distance	84	83	84	86	100	100	93	87	88	87	88	99	100	96	87	87	86	89	99	100	96
	Euclidean distance	88	87	88	87	94	93	100	91	91	91	90	96	96	100	90	90	90	90	96	96	100
		SWOW-RW							SWOW-RW-SVD							SWOW-PPMI-SVD						
Function2	Cosine similarity	100	82	83	79	77	76	79	100	89	90	83	81	81	81	100	96	99	94	89	89	89
	Dot product	82	100	98	94	93	93	95	89	100	97	91	88	87	89	96	100	98	93	86	87	87
	Pearson correlation	83	98	100	94	93	93	94	90	97	100	91	89	88	89	99	98	100	94	89	89	90
	Spearman correlation	79	94	94	100	93	93	91	83	91	91	100	89	88	88	94	93	94	100	89	89	87
	Augmented city block distance	77	93	93	93	100	100	96	81	88	89	89	100	99	96	89	86	89	89	100	99	96
	City block distance	76	93	93	93	100	100	96	81	87	88	88	99	100	95	89	87	89	89	99	100	96
	Euclidean distance	79	95	94	91	96	96	100	81	89	89	88	96	95	100	89	87	90	87	96	96	100
		Cosine similarity	Dot product	Pearson correlation	Spearman correlation	Augmented city block distance	City block distance	Euclidean distance	Cosine similarity	Dot product	Pearson correlation	Spearman correlation	Augmented city block distance	City block distance	Euclidean distance	Cosine similarity	Dot product	Pearson correlation	Spearman correlation	Augmented city block distance	City block distance	Euclidean distance
		Function1							Function1							Function1						

Fig. 6. Spearman correlations (multiplied by 100) between weighted similarity function predictions for all nine representations, averaged over all eight categories. Weighted functions are learned separately for each category.

3.2.2 Interim Discussion

Our primary finding from this section was that category-specific weighted functions of similarity could be learned via supervision from human similarity judgments between co-hyponyms of a category, suggesting that some dimensions matter more than others in the perception of lexical similarity. Out-of-sample predictions of similarity judgments from weighted functions over vector representations were greatly improved beyond their unweighted counterparts (on average, $r = .22$), with the best (representation, function) pair being unreduced SWOW-RW vectors with a city block distance function (average $r = .774$). The introduction of weights to similarity functions left the relative performance of different vector representations largely unchanged from the unweighted case, with vectors from association norms continuing to perform the best, and vectors specialized for similarity failing to outperform their unspecialized counterparts. More surprising was that the introduction of weights reduced the performance differences among similarity functions. This led to the finding that the introduction of weights to similarity functions caused these functions to make extremely similar predictions, suggesting identifiability issues among the weighted similarity functions we consider, *at least for the vector representations and stimuli we have explored*. We return to this issue in the General Discussion.

While identifying the precise weighted function eludes us for the moment, we do believe that, to the extent the weighted functions make better predictions than do unweighted functions, the weighted functions constitute better theories of similarity judgment among co-hyponyms. Thus far, however, we have built models predicting similarity by training and testing only on the similarity judgments for a particular category. Thus, the weights are potentially

specific only to a particular context or category, which would be consistent with a broad array of effects of context-dependence in judgment of similarity (Garner, 1962; Tversky, 1977; Bransford et al., 1989; Pomerantz & Lockhead, 1991; Goldstone et al., 1977). At the same time, Peterson et al. (2018) showed in the visual domain that weights trained on a mix of categories were able to generalize to an unseen category, thus demonstrating a great degree of context-independence, or category-generalizability, to the learned weights. The degree of context-(in)dependence or category specificity/generalizability of learned weighted similarity functions for the present data is thus somewhat *a priori* uncertain. In the next section, we turn to exploring this issue.

3.3 Category-general weighted functions

To better understand the category- or context-dependence of similarity judgment, we conduct two analyses. First, we attempt to fit models to seven categories at once, and predict judgments for pairs from the held-out eighth category. Second, we attempt to fit models on all eight categories at once, and predict judgments for unseen pairs from the trained categories (but again, unseen pairs in this case are still co-hyponyms). In the supplemental materials, we also attempt to learn weighted functions that predict similarity judgments in SimLex-999, which, as mentioned in the introduction, is a dataset of similarity judgments between pairs of words from a wide swath of the lexicon (e.g., good-bad, bed-hospital, cattle-beef). Attempting to model SimLex-999 thus constitutes another attempt to learn category-general weighted similarity functions, but for the sake of space, and because this analysis is largely similar to the first two, we do not explain it further here.

First, we attempt to fit models to similarity judgments from seven categories, and evaluate these models' predictions of similarity for an eighth, held-out category. As in the

previous section, we use non-negative L2-regularized regression, and conduct cross-validation on seven categories to select the strength of the regularization and find the dimension weights, and then evaluate predictions on the eighth category. Fig. 1C contains a heatmap summarizing performance averaged over all eight categories of words; Fig. 7 contains the same performance by category. It is apparent that, across vector types and similarity functions, performance is much worse compared to predictions generated by models trained and tested on only one category (compare Figs. 1B and 1C; mean decrease in performance of $r = -.24$) and is actually even slightly worse than predictions generated by unweighted models (compare Figs. 1A and 1C; mean decrease in performance of $r = -.02$). Further, no individual (representation, similarity function) pair outperforms the best unweighted (representation, similarity function) pairs (SWOW-RW-SVD with Spearman correlation, or SWOW-PPMI-SVD with cosine similarity or Pearson correlation, which both achieved $r = .61$). Overall, this pattern of results suggests that learning domain-general weights is generally not feasible *when the training and testing sets do not share categories*.

Fig. 7. Performance of every (representation, similarity function) pair in the hold-one-category-out setting. Cells indicate the Spearman correlation between mean human judgments and similarity function predictions, multiplied by 100 (for readability).

One explanation for the poor performance in the held-out setting is that category-general weights *could* be learned, but that they require at least some training data from every category. Hence, we next considered weighted similarity functions trained on pairs from all categories, and tested on unseen co-hyponym pairs from those same categories. Again, we use non-negative L2-regularized regression, and conduct nested cross-validation on word pairs randomly chosen from all eight categories to select the strength of the regularization and find the dimension weights, and then evaluate predictions on the remaining pairs from all eight categories. Because the distribution of human judgments varied by category (see Fig. S7 and Table S1 in the supplemental materials), such that people systematically judged the similarity of, e.g., vegetables ($M=3.08$, $SD=0.85$), to be higher than that of, e.g., sports ($M=2.06$, $SD=0.77$), it was important to compute correlations between predictions and actual judgments *by category*, and then average these correlations to get an overall measure of out-of-sample performance.¹¹

We visualize performance broken down by category, vector type, and similarity function in Fig. 8, and averaged over categories in Fig. 1D. Visual comparison to Fig. 1A makes clear that performance in the present setting is above that in the unweighted case (average increase of $r = .09$), but comparison to Fig. 1B makes clear that performance in the mixed training setting is below that in the within-category weighted case (average decrease of $r = .13$). We still see the same relative ordering of vector types – association norms perform best; specialized vectors fail to outperform their unspecialized counterparts; vectors trained on larger corpora perform better – as well as the highly similar performance of different weighted similarity functions. Further, in contrast to the unweighted function setting and the weighted

¹¹ If we evaluated correlations between predicted and actual judgments while ignoring category, we could induce a Simpson's Paradox that would make out-of-sample predictions appear spuriously accurate.

functions trained on a per-category basis, category-level split-half reliability now has only a very weak correlation with category average performance, [posterior median $r = .05$, 95% credible interval $(-.46, .55)$]. Yet, there is non-negligible variability in category predictability ($sd=.09$), with vehicles being predicted only at $r = .43$ (worse than the unweighted case, $r = .63$) and furniture being predicted at $r = .69$. What, then, accounts for the variability in category predictability in this mixed training and testing setting? One might have suspected that the number of pairs per category would be a factor, as larger categories (e.g. birds, with 435 pairs, cf. furniture with 190) would contribute more items to training and thus could push the weights towards their preferred region of weight space. However, category performance and number of pairs per category are, if anything, negatively correlated [posterior median $r = -.15$, 95% credible interval $(-.65, .36)$]. We tentatively suggest that category similarity to the *other* categories in the training set may be a factor: two of our eight categories, sports and professions, concern abstract concepts [both categories' concreteness ratings according to the Brysbaert et al. (2014) norms: $M=4.4$, $SD=.3$], while the rest concern concrete objects (means above 4.8, and SD 's below .24). Indeed, professions and sports have the lowest and third lowest performance in the mixed setting. Finally, in this setting, the best (representation, metric) pair is SWOW-RW with weighted Euclidean distance, which achieves an average performance of $r = .73$, which is about .04 points below the best (representation, function) pair when fitting weights on a per-category basis (SWOW-RW with city block, $r = .77$) and .12 points above the best unweighted (representation, function) pairs (SWOW-RW-SVD with Spearman correlation, and SWOW-PPMI-SVD with cosine or Pearson correlation, $r = .61$). Overall, when using training data from all categories to be tested, learning a single set of weights leads to performance above unweighted similarity

functions, but below weighted functions fit separately to each category, suggesting that *somewhat* domain-general weights can be learned. The supplemental materials contain an MDS-based analysis of the category-specific weight vectors that helps explain the difficulty of learning category-general weights: each category has a strong preference for weight vectors in idiosyncratic regions of weight space.

Fig. 8. Performance of every (representation, similarity function) pair in the mixed category setting. Cells indicate the Spearman correlation between mean human judgments and similarity function predictions, multiplied by 100 (for readability).

The above analyses suggest that different categories induce attention to different dimensions of a single latent space but say little about what information these dimensions might contain. To better understand this, we conducted a final analysis, using additional data from three prior studies: similarity judgments for a representative set of 60 personality traits (Richie et al., 2020), agreeableness scores derived from factor analysis for 449 traits including the 60 above (Ashton et al., 2015), and ratings of significance for 200 professions (i.e., how meaningful a job is perceived to be, Richie et al., 2019). Using non-negative L2-regularized regression, we fit trait agreeableness scores to the word2vec GoogleNews vectors for all 449 traits, and likewise for profession significance scores and all 200 professions, and extracted these two weight vectors. We then fit weighted dot product functions to trait similarity judgments and profession similarity judgments (again using word2vec GoogleNews vectors) and extracted these two weight vectors. Finally, we Pearson correlated the former two weight vectors with the latter two weight vectors. Table 3 has these correlations. As expected, the weight vector for judging trait similarity is more like that for judging trait agreeableness than for judging occupational significance, and vice versa (to a smaller extent) for the weight vector for judging profession similarity. This suggests that the dimensions people attend to when judging similarity of traits are closer to dimensions reflecting agreeableness than to those reflecting (occupational) significance. Conversely, when people judge similarity of professions, they attend to dimensions reflecting occupational significance slightly more than they attend to dimensions reflecting agreeableness. That the latter difference is smaller is perhaps sensible, since professions might also vary in stereotypical agreeableness (nurses and teachers may be perceived as similar partly in virtue of their stereotypical agreeableness). However, we caution that we have not quantified uncertainty in these

correlations¹², so this analysis must remain only suggestive for the time being. Still, we suspect that for all our categories, the learned weights reflect attention to interpretable dimensions of variation that are specific to a category. In the supplemental materials section *Analysis of feature types*, we conduct an exploratory analysis of the types of features common to our different categories, which may provide hints about the kind of information that category-specific weights emphasize.

Table 3. Correlations between (rows) learned weights for trait and profession similarity judgments, and (columns) learned weights for trait agreeableness and profession significance.

	Agreeableness weights	Significance weights
Trait similarity weights	0.18	0.04
Profession similarity weights	0.05	0.07

3.3.1 Interim Discussion

In this section, we attempted to learn weighted similarity functions that generalized across several categories, contexts, or domains. We attempted this in two ways: (1) fitting weighted functions to seven categories at once, and predicting judgments for pairs from the held-out eighth category, and (2) fitting weighted functions on all eight categories at once, and predicting judgments for unseen pairs from the trained categories. In an analysis reported in the supplemental materials, we also (3) fit and tested weighted functions on judgments from SimLex-999. In each case, learning category-general weights improved performance over unweighted functions to a much lesser degree [or not at all, in setting (1)] compared to fitting weights on a per-category basis (compare to previous section *Category-specific weighted similarity functions*). Preliminary analysis of the learned weight vectors for similarity among

¹² Each individual weight is not independent of the others, meaning we can not simply calculate confidence intervals around the correlations. The correlations function here more as similarity metrics.

professions or traits suggested why: category-specific weights reflected attention to dimensions of variation idiosyncratic to those categories (significance for professions and agreeableness for traits). In all, this pattern of results suggests that the dimensions of word vector representations that are emphasized during similarity judgment are largely dependent on context, which in the present case means the category(ies) or domain(s) from which words are drawn. Given the importance of this set of findings, we save further interpretation for the General Discussion.

For now, we only note that the augmented weighted city block distance function, borrowed from the Bayesian Analogy with Relational Transformations (Lu et al., 2012, 2019; Chen et al., 2017), did not outperform other weighted similarity functions in the present domain-general settings. This is somewhat unexpected given the success of this function in learning domain-general weights for relations like X-ISA-Y or X-CAUSES-Y, or, indeed, X-SYNONYM-Y (Lu et al., 2019). Why does BART succeed with these relations, and specifically with predicting synonymy, regardless of domain or category, while the same similarity function can not accurately predict similarity ratings of co-hyponym pairs in our settings? We suspect that our data simply present a more challenging task. That is, our data require distinguishing many levels of similarity among items that are highly similar to each other, relative to random words in the lexicon (e.g., airplane and skateboard are dissimilar in the context of other pairs of vehicles, but highly similar in the context of the entire lexicon; see also the Tversky 1977 study of similarity within South American or European countries, discussed in the introduction).

4. General Discussion

In the current work, we evaluated and compared seven similarity functions provided with nine types of word vectors derived from text and association norms, on their ability to

predict pairwise similarity judgments among co-hyponyms (e.g., sofa, bed) of 8 different categories (e.g. furniture). We obtained several key results:

First, among unweighted similarity functions, cosine similarity, despite its popularity in applications with vector space models, did not show general or substantial advantages over city block distance, Euclidean distance, or dot product. Further, cosine performed about as well as Pearson correlation, which has little precedence in applications involving word vectors. In fact vectors derived from free association norms and reduced with SVD, combined with *Spearman* correlation as a similarity function tied for the best (representation, unweighted function) pair ($r = .61$), and performed .04 points better than the next best (representation, unweighted function) pair (SWOW-RW, Pearson correlation), and .08 points better than the best (text-based vector, function) pair (fastText-600B-CC, Pearson correlation).

Second, we found no advantage for text-based vectors specialized for similarity over their unspecialized counterparts (unlike on the similarity benchmark SimLex-999), and found that word vectors derived from free association norms generally achieve small but reliable advantages ($\sim r = .05$) over text-based vectors, when using either unweighted or weighted similarity functions.

Third, using human similarity judgments to supervise learning of category-specific weights on dimensions with regularized non-negative regression, we substantially improved prediction of similarity judgments beyond unweighted similarity functions (average improvement of $r = .22$) and found little difference between weighted similarity functions. Comparison of the predictions of different weighted similarity functions and a model recovery

exercise both suggested that each weighted similarity function makes very similar predictions to the others, raising questions about the identifiability of these weighted functions.

Fourth, in both the unweighted and especially the weighted function settings, category-level split-half reliability accounted for a moderate or large amount of variability in category predictability, but additional contributors to variability in category predictability (e.g., frequency and familiarity) are unclear and are therefore potential foci of future research.

Finally, and perhaps most importantly, we found that learning category-general or context-independent weights for *lexical* similarity judgment is generally difficult. As we attempt to learn weights in increasingly general settings – (1) learning separate weights for each category, (2) learning a single weight vector for all of our eight categories simultaneously, and (3) learning a single weight vector from seven categories and applying it to an eighth – we found that the benefit of weights weakens more and more. The introduction of weights actually became slightly harmful when attempting to learn weights on seven categories and generalizing those weights to an eighth, held-out category. An examination of the category-specific weights (reported in the supplemental materials) yielded some insight into this finding: different categories had a strong tendency for weights in distinct, non-overlapping regions of weight space. A preliminary analysis suggested that, in the case of professions, these weights appear to emphasize dimensions corresponding to occupational significance, but for personality traits, dimensions corresponding to agreeableness. This final set of findings leads us to conclude that the dimensions of word vector representations that are emphasized during similarity judgment are largely dependent on context, which in the present case means the category(ies) or domain(s) from which words are drawn.

There are several ways in which this work extends earlier studies comparing various semantic representations' and similarity functions' ability to model semantic behavior, particularly similarity and relatedness judgments (Landauer & Dumais, 1997; Burgess & Lund, 2000; Rohde et al., 2006; Bullinaria & Levy, 2007, 2012, Mander et al., 2017; Pereira et al., 2016; Hill et al., 2014; De Deyne et al., 2015; Hill et al., 2015; Wieting et al., 2015; Gerz et al., 2016; Ponti et al., 2020). First, our similarity judgments among hundreds of co-hyponym pairs in eight different categories place much more challenging demands on representations and similarity functions than datasets used in the past, like SimLex-999 (Hill et al., 2015) or the synonym portion of the TOEFL (Landauer & Dumais, 1997). As we argued earlier, these latter datasets do not require distinguishing the many levels (potentially hundreds) of perceived similarity among pairs of items within a category (e.g. fruits) that are generally quite similar to each other (at least relative to other items in the lexicon). Yet it is the former, within-category similarity judgments that cognitive scientists are often interested in, for modeling, e.g., categorization (Nosofsky, 1984, 1986). Second, we compare very recent vector representations from free association norms (De Deyne et al., 2019) as well as from text, including text-based vector representations specialized for similarity, which perform well on SimLex-999 (Wieting et al. 2015; Ponti et al., 2020). Third, our work is the first to examine (a) the improvements in accuracy achieved by weighted similarity functions operating over dense, high-dimensional vector representations of words, (b) how these weights depend on the context – here, category of items – of similarity judgment, and (c) the effect of weights on the predictions made by different types of similarity functions, and on the identifiability of these functions. In the sections that follow, we further discuss implications of our research.

4.1 Representations

Across our studies of co-hyponym similarity judgments, at least four general findings regarding the representations that best explained similarity judgments recurred. First, and most obviously, text-based representations benefited from training on larger corpora. As discussed in the first *Interim Discussion*, this result is straightforward: larger corpora provide more contexts of occurrence for judgment target words, and hence more information with which word vectors are specified, which may be especially important for the present fine-grained similarity judgments. Second, for text-based vectors trained on roughly equally sized corpora, the training algorithm – fastText or GloVe – seemed to matter little, which is potentially consistent with previous work showing more data trumps smarter algorithms (Recchia & Jones, 2009), and that there are minimal differences between prediction and count models when model parameters are tuned carefully (Levy et al., 2015). Our last two findings are potentially more novel, and we take more time to carefully consider them.

First, in all four weight training settings visualized in Figs. 1, 2, 5, 7, and 9 – Unweighted, Per-category weights, Hold one category out, and Mix of categories – vectors from association norms performed slightly better than the best vectors from text, with advantages ranging from $r = .05$ to $r = .11$. Notably, this superiority held even when comparing SWOW-PPMI-SVD to our word2vec representations, whose SGNS learning algorithm is equivalent to factorization (e.g., SVD) of a PMI matrix (Levy & Goldberg, 2014). As we noted in the first interim discussion, this result is consistent with other work comparing text-based vectors and vectors from association norms (De Deyne et al., 2015; Vankrunkelsven et al., 2018; De Deyne et al., 2021), but we have extended these comparisons by considering more text-based

representations and similarity functions. Like De Deyne et al. (2015) and Vanrunkelsven et al. (2018), we suggest that free association data contain different information from text corpora, information that is apparently more relevant to predicting similarity (and other semantic phenomena). For example, De Deyne et al. (2015) suggest that pragmatic constraints in communication, like the need to be informative, mean that people do not often state obvious facts like “Bananas are yellow” in text corpora, even though yellowness is a salient property of bananas. In contrast, in free association, people can and do respond ‘yellow’ when cued with ‘banana’. Of course, to accurately model similarity judgments, a vector representation must have information reflecting such central, obvious properties, probably even more so than less obvious, more pragmatically felicitous facts like “Bananas are rich in potassium”, which might be better reflected in text.

One might object, however, that vectors based on free association norms ‘build in’ the very structure they are attempting to explain. A similar critique has been leveled against use of semantic networks based on free association data to model behavior in category recall tasks (“Name all the animals you can think of in a minute”; Abbot et al., 2015; Jones et al., 2015). Hills et al. (2012) showed that category recall behavior had signatures of optimal foraging in memory that resembled search (e.g. for food) in physical environments. They then showed that these data were well-explained by a two-stage cue-switching model applied to a semantic space learned from text corpora. Abbot et al. (2015) then showed that a random walk on a semantic network constructed from free association data could reproduce these same patterns. However, as Jones et al. (2015) note, Abbot et al.’s (2015) process model uses human behavioral data from a task that is very similar to the behavior they attempt to explain, so it is not clear that the Abbot et

al. model has any explanatory power. We are sympathetic to this concern (although we feel free association is much more like category recall than similarity judgment), and therefore refrain from claiming that vectors from free association constitute psychological theories of semantic representation – let alone theories of learning – the same way that distributional models based on text can be. Nonetheless, the superiority of the association-based vectors over text-based vectors does suggest the latter are missing predictable variation in similarity (possibly for the reasons discussed in the last paragraph), and thus the results from the association-based vectors are psychologically relevant – they clarify how an account of learning from just text is insufficient. Finally, even if the association-based vectors do not constitute a psychological theory of representation, our work shows they can be convenient proxies of similarity, which might be practically useful. We discuss this further in the upcoming section *Implications for Cognitive Modeling*.

Second, we found that, in contrast to results with the similarity benchmark SimLex-999, where similarity (synonymy) is distinguished from relatedness and antonymy, we found no advantage for word vectors specialized for similarity, in any of our four weight training settings summarized in Fig. 1. In fact, we often observed that these specialization techniques *harmed* prediction of similarity judgments, especially in the per-category weights setting. It is not clear why these techniques are unhelpful in the present setting, but as mentioned in the introduction, the answer may have to do with the goals of these techniques. That is, these techniques are largely focused on disentangling relatedness from similarity, and pushing the vectors for antonyms apart and the vectors for synonyms together. Conflating relatedness with similarity was unlikely to be a problem with co-hyponyms (related but arguably dissimilar co-hyponyms

like ‘burger’ and ‘fries’ are perhaps rare), nor are there any antonym pairs in our categories (this likely would have been an issue if our categories included, e.g., traits or emotions), and very few of the pairs in our dataset are exactly synonymous, even though they may be very similar (e.g., sweatshirt and sweater). It may also be the case that there are simply not enough words relevant to our categories in the lexical constraints that are used to specialize word vectors. Still, it is perhaps puzzling that these specialization techniques should *hurt* performance, except to the extent that the specialization techniques move vectors from already (near-)optimal positions in space.

Of course, there is no reason that people must be using exactly one kind of representation when making similarity judgments. It could be the case that similarity judgment is best modeled as some operation over representations that are combinations of those from text, free association norms, and other sources. Moreover, for certain domains, featural representations are likely to be required. Numbers, for example, can vary both in magnitude, which is continuous (like the dimensions of the vector representations we use), but also in whether they are prime numbers, which is a discrete property (Navarro and Lee, 2003). A clear direction for future work is to examine the power of combinations of representations, be they featural representations or spatial-like representations from text or free association norms, to model perception of lexical similarity.

4.2 Similarity functions

One of the major goals of this work was to understand the similarity functions that best describe people’s judgments of similarity, i.e., to discriminate among different computational-level theories of similarity judgment (Marr, 1982). However, in contrast to the

substantial variation we found in representations' ability to predict similarity judgments in various settings, we found much less variation among similarity functions. This was true in our unweighted setting, where Pearson correlation and cosine similarity had only small advantages (difference in r less than .06) over other functions, but especially in our per-category weighted setting, where all weighted similarity functions achieved average performance between $r = .67$ and $r = .69$. This is perhaps somewhat unexpected given work briefly mentioned in section 1.2.2, showing that common features models like (weighted) cosine, dot product, and correlation better describe human similarity judgments in linguistic stimuli than do distinctive features models like (augmented) city block and Euclidean distance (Gati & Tversky, 1984; Ritov, Gati, & Tversky, 1990; Goldstone & Son, 2004).

We take different views of this for unweighted and weighted similarity functions. To the extent that the unweighted similarity functions make worse predictions of similarity judgments, we take the unweighted similarity functions less seriously as cognitive models of similarity judgment, and more as convenient approximations – in that they need no supervision to learn weights – to similarity judgments (see upcoming section *Implications for Cognitive Modeling*). Thus, the fact that different unweighted similarity functions make similarly good predictions – when averaged over representations and categories – is somewhat immaterial. We only reiterate what we have said previously, that the general dominance of cosine similarity in previous work is somewhat undeserved: Pearson correlation performed just as well overall, and cosine similarity was actually somewhat *worse* than Spearman correlation for SWOW-RW-SVD vectors (difference of $r = .06$); again, this latter (representation, similarity function) pair tied for the best overall performance among unweighted similarity functions applied to our eight

categories. This last finding is rather unanticipated, and we do not currently have a good explanation of why Spearman correlation pairs so well with SWOW-RW-SVD vectors.

However, to the extent that *weighted* similarity functions, at least for the per-category weighted and mixed category settings, make *better* predictions of similarity judgments than do unweighted functions, we *do* take weighted similarity functions seriously as cognitive models of similarity judgment, and thus we do have a more serious interest in identifying what weighted similarity function best describes human similarity judgments. On the basis of our results, we can suggest that human computation of similarity may involve a high-dimensional representational space common to many different categories and items (indeed, the thousands or millions of words in typical text-based vector spaces), but with weighted similarity functions whose weights vary somewhat by category or context of similarity judgment (we further discuss context effects in the final section). However, the precise functional form, e.g., weighted cosine or weighted city block distance – or more generally, common or distinctive features models (Navarro & Lee, 2004) – is difficult to determine from our results, owing to the substantial problems with identifiability among the weighted similarity functions we consider, *at least for the vector representations and stimuli we have explored*. We emphasize this point, as it is possible to conceive of pairs of stimuli that distinguish the present weighted similarity functions. For example, consider the weighted dot product and weighted city block or Euclidean distance functions. For the weighted dot product, a pair of words that has a similarly high value on every dimension (e.g., ~ 3) will be more similar than a pair of words that has a similarly low value on every dimension (e.g., ~ 1). These pairs will have identical (dis)similarity for weighted city-block or Euclidean distance. This is true as long there is at least one non-zero weight. Of course, it

seems that our pairs of words do not contain enough of these disambiguating pairs of stimuli to distinguish similarity functions effectively. Future work might therefore consider utilizing optimal experimental design tools to select stimuli that can distinguish these weighted functions (e.g., Myung & Pitt, 2009).

4.3 Implications for cognitive modeling

Estimates of similarity, or representations derived from similarity, are required for the operation of many cognitive models of phenomena like generalization and categorization (Shepard, 1987; Nosofsky, 1984, 1986, 2014). Ordinarily, a researcher interested in applying these models would collect similarity judgments between the relevant items (usually with pairwise Likert scale ratings, as we do here), and then use these measurements or derived representations in their model.¹³ Unfortunately, collecting similarity judgments can be a costly endeavor, whether one uses pairwise Likert scale ratings or more recent methods like the Spatial Arrangement Method (SpAM; Goldstone, 1994; Hout et al., 2013; Richie et al., 2020; but for concerns about SpAM’s validity, see Verheyen et al., 2016) and generalizations of the odd-one-out or triad task (e.g., Roads & Mozer, 2019).

Estimating similarity from off-the-shelf general purposes representations from text or free association norms may thus be useful. Our best-performing unweighted similarity models, SWOW-RW-SVD with Spearman correlation as a similarity function and SWOW-PPMI-SVD with cosine or Pearson correlation as similarity metrics, achieved an average Spearman

¹³ One might feel that we are implicitly endorsing the assumption that when similarities are needed for cognitive modeling, it entails a two step process: first deriving a representation from similarities, and then modeling a task like categorization. We acknowledge that this need not be the case. One can model data from a similarity rating task and a task like categorization, simultaneously (Lee, 2011). In our case, this could entail learning a single set of weights that reflects attention to latent dimensions in similarity rating, and categorization. We thank a reviewer for pointing this out.

correlation with human pairwise Likert scale judgments of $r = .61$. When a researcher requires similarity judgments between co-hyponyms, can not afford to collect them from human participants, but an approximation is acceptable, we can recommend using these (representation, similarity function) pairs. Of course, even SWOW is not limitless in its vocabulary, as it has under 13,000 words. Still, it contained all the words we selected for our studies, and we suspect it will be adequate for many researchers. For researchers that require similarity judgments for more obscure words, we recommend using text-based vectors that are trained on the largest possible corpora available, e.g., fastText-600B-CommonCrawl with cosine or Pearson correlation as the similarity functions, which achieved Spearman correlations with human judgments of $r = .53$. We also tentatively suggest researchers use ordinary word vectors that have not been further specialized according to methods like those of Wieting et al. (2015) or Ponti et al. (2020), as the excellent performance of these procedures on SimLex-999 appears not to extend to fine grained judgments of similarity among co-hyponyms. However, there are other methods of specialization that we have not tested (e.g., Mrkšić et al., 2017), and newer techniques are sure to emerge, so we are somewhat hesitant to generalize to the entire class of existing specialization procedures.

Further, if the number of items is high (~ 100) such that the number of pairs is exceedingly high (~ 5000), as might happen if the category is large (e.g., personality traits, foods, animals), then our per-category weighted functions suggest that it may be possible to collect just a sample of similarity judgments from among the possible pairs (e.g., 200-400 pairs), learn a weighted function for this category, and use this function to make highly accurate predictions among the remaining pairs. In this setting, we again recommend using vectors from association

norms, specifically the unreduced SWOW-RW norms. The weighted city block distance function very slightly outperformed other weighted functions with this representation, but of course our results suggest that many weighted functions will generally make equally good results.

Finally, while we did not see especially large differences between similarity functions, in both unweighted but especially weighted settings, if future work can identify the (weighted) similarity function(s) that best describe human similarity judgments, then it may be wise to use these functions in cognitive models. For example, the generalized context model, as typically implemented, assumes that categorization depends on the similarity between exemplars as computed by a weighted Minkowski distance passed through an exponential decay function. However, if it turned out that, say, weighted dot product actually better accounted for human similarity judgments among co-hyponyms, then perhaps this function ought to be used in the GCM, at least when learning novel categories over words. Indeed, Nosofsky et al. (2018), in an application of the GCM to modeling category learning for a set of naturalistic, visually complex rocks varying on 8 or 9 dimensions, made a similar suggestion, that “more sophisticated processing models may be needed for characterizing similarity relations among complex objects such as rocks” (pg. 20). While Nosofsky et al. had in mind a particular featural model of similarity we have not considered here – Tversky’s (1977) feature-contrast model – the point is that rich, naturalistic domains (be they visual or lexical/conceptual) may necessitate different similarity functions than were developed for modeling behaviors involving simple, artificial stimuli.

4.4 Implications for Natural Language Processing

We found that the present similarity judgments of co-hyponyms were more difficult to model with *unweighted* similarity functions over word vector representations than are predecessor similarity benchmarks like SimLex-999: Our best unweighted models achieved a mean correlation with human judgments of $r = .61$ (SWOW-RW-SVD with Spearman correlation as the similarity function, and SWOW-PPMI-SVD with cosine or Pearson correlation as the similarity function), while the best performance on SimLex-999 to date is $r = .79$ (obtained by GloVe-6B-Giga+Wiki, specialized with Ponti et al.’s. 2020 adversarial postspecialization method, using any similarity metric except (augmented) city block distance). To the extent that the present data capture important aspects of the perception of similarity not tapped by SimLex-999 and related datasets, the present similarity judgments may be useful as an additional benchmark for the intrinsic evaluation of new word vector representations (Hill et al., 2015; Gerz et al., 2016; Vulić et al., 2020). Just as word vectors that perform better on SimLex-999 tend to perform better on NLP tasks like lexical text simplification and dialogue state tracking (Ponti et al., 2020; Mrkšić et al., 2017), we suspect that representations that perform better on the present kinds of similarity data will perform better on a variety of natural language processing and artificial intelligence tasks, including the recommendation task described in the opening paragraph of this paper [but see Schnabel et al. (2015) and Tsvetkov et al. (2015) for NLP tasks where model performance on similarity rating is a poor predictor]).

At the same time, the context-sensitivity of the present kinds of judgments suggests that it may not be possible to derive a text-based representation, with or without specialization techniques, that perfectly respects similarity relations in all possible categories at once, even among fairly coherent, reasonably ‘natural’ categories like the ones considered here, *at least*

when combined with unweighted similarity functions. That is, in order to accurately predict human similarity judgments among co-hyponyms of a category from a single representational space that covers the millions of words and phrases in the lexicon, the same space needs to be stretched or compressed in ways that are largely idiosyncratic to a category or context. To the extent that this is true, learning a vector specialization function that works for every category seems unlikely to be possible given (a) it is difficult to adequately sample the extremely large, potentially infinite number of possible categories of words, and (b) the same item can participate in multiple categories or contexts, each of which may induce different dimensional weights. Thus, we somewhat caution against attempting to optimize performance of vector specialization techniques on the present kinds of data, as such an attempt may overfit.

4.5 Context effects in similarity judgment

Our final set of results was that similarity judgments among co-hyponyms were not well-modeled by weighted similarity functions that generalized across categories, contexts, or domains. In particular, we found that training weights on seven categories and testing these weights on an eighth category yielded performance no better than unweighted similarity functions. This latter finding contrasts sharply with closely related work by Peterson et al. (2018) who performed similar analyses with judgments of image similarity and image representations extracted from various layers (optimally, later layers, which encode more abstract information) of deep neural networks trained to perform image classification. They found that, when training on five categories and testing on a held-out sixth category, the proportion of variance explained was roughly double the variance explained by the unweighted dot product, for every held-out category (although Peterson et al. did find that training and testing on only single category still

achieved higher accuracy, suggesting even their visual stimuli induced somewhat category-specific dimensional weights). Thus, our inability to learn domain-general or context-independent weighted similarity functions is more in line with the large psychological literature on context effects in similarity judgment (e.g., Tversky, 1977; Barsalou, 1982; Medin et al., 1993; Goldstone et al., 1997; Close et al., 2009). Rather than similarity depending on some all-purpose function of distance on a static space, we found that the computation of similarity among words depends on the set of words given in a context, in ways that can be formally specified by supervised learning of dimensional weights.

What accounts for the difference between our results and those of Peterson et al. (2018)? We speculated in the introduction that this is largely a difference between words and images. In images, a relatively constant set of features or dimensions may be relevant to similarity, like color, shape, or texture, but for lexicalized, often more abstract concepts, different lexical categories seem likely to have different dimensions of variation. For example, professions vary in their perceived significance, while traits vary in their agreeableness. This suggestion is consistent with the classic *extension effect* mentioned in the introduction, according to which features influence similarity judgments more when they vary within an entire set of stimuli (Garner, 1962; Bransford et al., 1989; Pomerantz & Lockhead, 1991). In our final analysis of section 3.3, we provided preliminary evidence in support of this suggestion: the learned weight vectors for judgments of similarity corresponded to dimensions representing trait agreeableness, while the weight vectors for judgments of similarity among professions corresponded to dimensions representing perceived significance.

Additional principles may explain the context-dependence of weighting in similarity judgment. In the present work, we have focused on real-valued vector representations derived from text or association norms, but as we noted earlier, binary featural representations are another popular representational format (McRae et al., 2005; Buchanan et al., 2019). In this latter representational tradition, there exists the *size principle*. Assuming a common features model of similarity (recall, the weighted dot product), the size principle states that the weight on a feature is inversely proportional to the number of items possessing that feature in a given context (Tenenbaum & Griffiths, 2001; Navarro & Perfors, 2010). Indeed, this principle has been empirically validated for similarity judgments resembling our own, with the feature norms we describe earlier (Peterson & Griffiths, 2017; McRae et al., 2005). It is not clear that the size principle can be applied to the real-valued representations we tested in the present study, but there do exist binary-valued vector representations derived from text (e.g. Tissier et al., 2018), and it may be possible to test *both* the size principle and the extension effect on our judgments using such representations (in principle feature norms could serve this purpose, but again, many of our items are missing from the largest available feature norms, Buchanan et al., 2019). Interestingly, the size principle and the extension effect may be somewhat at odds. The perhaps simplest application of the extension effect to the present setting would predict the strongest weights on dimensions with the greatest variability, which in a featural representation would be a feature where half of the given items have the feature and half don't. The size principle, however, predicts a weight is strongest when exactly one item has a feature. It may thus be worthwhile to directly compare the ability of these two principles to explain the variation of weights in context-dependent similarity judgment among lexical stimuli, using the sorts of

representations we used here. It may also be that explaining the context-dependence of weights in lexical similarity judgment will require the development of entirely new principles and theories. In any case, this is perhaps the direction for future research that we consider most exciting: developing and testing theories, ideally formally specified, of how and why context shifts the use of a single representational space – embodied by text or association norm vectors – in naturalistic lexical similarity judgment. As Gershman (2017) said, “[d]espite its ubiquity, the origin of context-sensitivity [in similarity judgment] is still poorly understood—we lack a computational account of when and why context effects occur.” If such a formal theory could be produced, it should be possible to predict the weights for a given set of items across different categories, and hence make highly accurate predictions of lexical similarity, without having to collect any human judgments to supervise learning of the weights. Our data – hosted at https://osf.io/d7fm2/?view_only=c5ba5d34a5e34ff3970a652c07aad5c – may be useful to other researchers in developing such theories.

Acknowledgments

Thanks to members of the Computational Behavioral Science Lab, attendees of SemSpace 2020, and attendees of the 60th Annual Meeting of the Psychonomic Society for their feedback. This work was funded by National Science Foundation grant SES-1626825, and the Alfred P. Sloan Foundation.

References

- Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, 122, 558-569.
- Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001, January). On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory* (pp. 420-434). Springer, Berlin, Heidelberg.
- Ashton, M. C., Lee, K., & Boies, K. (2015). One-through six-component solutions from ratings on familiar English personality-descriptive adjectives. *Journal of Individual Differences*, 36(3), 183.
- Barsalou, L. W (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, 10, 82-93.
- Bhatia, S., Richie, R., & Zou, W. (2019). Distributed semantic representations for modeling human judgment. *Current Opinion in Behavioral Sciences*, 29, 31-36.
- Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *The Journal of General Psychology*, 30, 149 –165.
<http://dx.doi.org/10.1080/00221309.1944.10544467>
- Bransford, J. D., Franks, J. J., Vye, N. J., & Sherwood, R. D. (1989). New approaches to instruction: Because wisdom can't be told. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 470-497). Cambridge: Cambridge University Press.
- Bruni, E., Boleda, G., Baroni, M., & Tran, N. K. (2012). Distributional semantics in technicolor. In *Proceedings of the 50th Annual Meeting of the Association for Computational*

- Linguistics: Long Papers-Volume 1* (pp. 136-145). Association for Computational Linguistics.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904-911.
- Buchanan, E. M., Valentine, K. D., & Maxwell, N. P. (2019). English semantic feature production norms: An extended database of 4436 concepts. *Behavior Research Methods*. 1-15. <https://doi.org/10.3758/s13428-019-01243-z>
- Bullinaria, J. A., & Levy, J. P. (2007). Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior research methods*, 39(3), 510-526.
- Bullinaria, J.A. & Levy, J.P. (2012). Extracting Semantic Representations from Word Co-occurrence Statistics: Stop-lists, Stemming and SVD. *Behavior Research Methods*, 44, 890-907.
- Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory. *Cognitive dynamics: Conceptual and representational change in humans and machines*, 13, 17-56.
- Cawley, G. C., & Talbot, N. L. (2010). On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11, 2079-2107.
- Chen, D., Lu, H., & Holyoak, K. J. (2017). Generative inferences based on learned relations. *Cognitive science*, 41, 1062-1092.
- Close, J., Hahn, U., & Honey, R. C. (2009). Contextual modulation of stimulus generalization in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 35(4), 509.

- Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, 23(3), 371-414.
- De Deyne, S., Navarro, D. J., & Storms, G. (2013). Better explanations of lexical and semantic cognition using networks derived from continued rather than single word associations. *Behavior Research Methods*, 45, 480–498.
- De Deyne, S., Navarro, D. J., Collell, G., & Perfors, A. (2021). Visual and Affective Multimodal Models of Word Meaning in Language and Mind. *Cognitive Science*, 45(1), e12922.
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior research methods*, 51(3), 987-1006.
- De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M. J., Voorspoels, W., & Storms, G. (2008). Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior research methods*, 40(4), 1030-1048.
- De Deyne, S., Verheyen, S., & Storms, G. (2015). The role of corpus size and syntax in deriving lexico-semantic representations for a wide range of concepts. *The Quarterly Journal of Experimental Psychology*, 68(8), 1643-1664.
- Garner, W. R. (1962). *Uncertainty and structure as psychological concepts*. New York: Wiley.
- Garner, W. R. (1974). *The Processing of Information and Structure*. Potomac, MD: Erlbaum.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. *Cognitive Psychology*, 16(3), 341-370.
- Gershman, S. (2017). Similarity as inference. Unpublished manuscript. Retrieved from <https://pdfs.semanticscholar.org/0ad0/8f6d2daed2446652c0e9eb9b79ac7e684e99.pdf>

- Gerz, D., Vulić, I., Hill, F., Reichart, R., & Korhonen, A. (2016). SimVerb-3500: A Large-Scale Evaluation Set of Verb Similarity. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 2173-2182).
- Glavaš, G., Ponti, E. M., & Vulić, I. (2019,). Semantic Specialization of Distributional Word Vectors. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): Tutorial Abstracts*.
- Goldstone, R. (1994). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers*, 26, 381–386. <https://doi.org/10.3758/BF03204653>
- Goldstone, R. L., Medin, D. L., & Halberstadt, J. (1997). Similarity in context. *Memory & Cognition*, 25(2), 237-255.
- Goldstone, R. L., & Son, J. Y. (2005). Similarity. In K. J. Holyoak, & R. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 13–36). Cambridge, England: Cambridge University Press.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2), 211.
- Gruenewald, P. J., & Lockhead, G. R. (1980). The free recall of category examples. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 225–240.
<http://dx.doi.org/10.1037/0278-7393.6.3.225>
- Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 14(6), 1006-1033.

- Hahn, U. (2014). Similarity. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(3), 271-280.
- Hill, F., Cho, K., Jean, S., Devin, C., & Bengio, Y. (2014). Embedding word similarity with neural machine translation. *arXiv preprint arXiv:1412.6448*.
- Hill, F., Reichart, R., & Korhonen, A. (2015). SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4), 665–695.
- Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological review*, 119(2), 431.
- Hout, M. C., Goldinger, S. D., & Ferguson, R. W. (2013). The versatility of SpAM: A fast, efficient, spatial method of data collection for multidimensional scaling. *Journal of Experimental Psychology: General*, 142(1), 256–281. <https://doi.org/10.1037/a0028860>
- Hout, M. C., Papesh, M. H., & Goldinger, S. D. (2013). Multidimensional scaling. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(1), 93-103.
- Howard, M. W., & Kahana, M. J. (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46, 85–98. <http://dx.doi.org/10.1006/jmla.2001.2798>
- Johnson, M. D., Lehmann, D. R., & Horne, D. R. (1990). The effects of fatigue on judgments of interproduct similarity. *International Journal of Research in Marketing*, 7, 35– 43.
doi:10.1016/0167-8116(90)90030-Q
- Jones, M. N., Hills, T. T., & Todd, P. M. (2015). Hidden processes in structural representations: A reply to Abbott, Austerweil, and Griffiths (2015). *Psychological Review*, 122, 570–574.
DOI: <https://doi.org/10.1037/a0039248>
- Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1), 1.

- Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. *Oxford handbook of mathematical and computational psychology*, 232-254.
- Jurafsky, D. & Martin, J.H. (in prep). Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition. (3rd edition.) Prentice-Hall.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430–454. [https://doi.org/10.1016/0010-0285\(72\)90016-3](https://doi.org/10.1016/0010-0285(72)90016-3)
- Keeney, R. L., Raiffa, H., & Meyer, R. F. (1993). *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.
- Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29(1), 1-27.
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299-321.
- 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. doi:10.1037/0033-295X.104.2.211
- Lee, M.D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55, 1-7.
- Lenci, A. (2018). Distributional models of word meaning. *Annual review of Linguistics*, 4, 151-171.
- Levy, O., & Goldberg, Y. (2014). Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems* (pp. 2177-2185).

- Levy, O., Goldberg, Y., & Dagan, I. (2015). Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3, 211-225.
- Lu, H., Chen, D., & Holyoak, K. J. (2012). Bayesian analogy with relational transformations. *Psychological Review*, 119(3), 617–648. <https://doi.org/10.1037/a0028719>
- Lu, H., Wu, Y. N., & Holyoak, K. J. (2019). Emergence of analogy from relation learning. *Proceedings of the National Academy of Sciences*, 116(10), 4176-4181.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavioral Research Methods, Instrumentation, and Computers*, 28, 203-208.
- Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. San Francisco: W. H. Freeman and Company.
- 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(4), 547–559. <https://doi.org/10.3758/BF03192726>
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological review*, 100(2), 254.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111–3119).

- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., & Joulin, A. (2018). Advances in pre-training distributed word representations. *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*.
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes Factors for common designs*. R package version 0.9.12-4.2. <https://CRAN.R-project.org/package=BayesFactor>
- Mrkšić, N., Vulić, I., Séaghdha, D. Ó., Leviant, I., Reichart, R., Gašić, M., ... & Young, S. (2017). Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the association for Computational Linguistics*, 5, 309-324.
- Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological review*, 116(3), 499.
- Navarro, D. J. and Lee, M. D. (2003). Combining dimensions and features in similarity-based representations. In S. Becker, S. Thrun and K. Obermayer (Ed.) *Advances in Neural Information Processing Systems* (pp. 67-74).
- Navarro, D. J., & Lee, M. D. (2004). Common and distinctive features in stimulus similarity: A modified version of the contrast model. *Psychonomic Bulletin & Review*, 11(6), 961-974.
- Navarro, D. J., & Perfors, A. F. (2010). Similarity, feature discovery, and the size principle. *Acta Psychologica*, 133(3), 256–268.
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36, 402–407. doi:10.3758/BF03195588
- Newman, M. E. J. (2010). *Networks: An introduction*. Oxford: Oxford University Press, Inc.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114.

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.

Nosofsky, R.M. (2014). The generalized context model: an exemplar model of classification. In M. Pothos and A. Wills (Eds.), *Formal Approaches in Categorization*. Cambridge University Press: New York.

Nosofsky, R. M., Sanders, C. A., & McDaniel, M. A. (2018). Tests of an exemplar-memory model of classification learning in a high-dimensional natural-science category domain. *Journal of Experimental Psychology: General*, 147(3), 328.

O'Connor, C. M., Cree, G. S., & McRae, K. (2009). Conceptual hierarchies in a flat attractor network: Dynamics of learning and computations. *Cognitive science*, 33(4), 665-708.

Patel, A., Sands, A., Callison-Burch, C., & Apidianaki, M. (2018). Magnitude: A fast, efficient universal vector embedding utility package. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2018*, 120–126.

Pedregosa F., et al. (2011) Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.

Pereira, F., Gershman, S., Ritter, S., & Botvinick, M. (2016). A comparative evaluation of off-the-shelf distributed semantic representations for modelling behavioural data. *Cognitive neuropsychology*, 33(3-4), 175-190.

- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- Peterson, J. C., Abbott, J. T., & Griffiths, T. L. (2018). Evaluating (and Improving) the Correspondence Between Deep Neural Networks and Human Representations. *Cognitive Science*, 42(8), 2648–2669. <https://doi.org/10.1111/cogs.12670>
- Peterson, J. C., & Griffiths, T. L. (2017). Evidence for the size principle in semantic and perceptual domains. *arXiv preprint arXiv:1705.03260*.
- Pomerantz, J. R., & Lockhead, G. R. (1991). Perception of structure: An overview. In G. R. Lockhead & J. R. Pomerantz (Eds.), *The perception of structure* (pp. 2-23). Washington: American Psychological Association.
- Ponti, E., Vulić, I., Glavaš, G., Mrkšić, N., & Korhonen, A. (2020). Adversarial propagation and zero-shot cross-lingual transfer of word vector specialization. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, 282-293. <https://doi.org/10.17863/CAM.35175>
- Recchia, G., & Jones, M. N. (2009). More data trumps smarter algorithms: Comparing pointwise mutual information with latent semantic analysis. *Behavior research methods*, 41(3), 647-656.
- Richie, R., White, B., Bhatia, S., & Hout, M. C. (2020). The spatial arrangement method of measuring similarity can capture high-dimensional semantic structures. *Behavior Research Methods*, 52, 1906–1928.

- Richie, R., Zou, W., & Bhatia, S. (2019). Predicting high-level human judgment across diverse behavioral domains. *Collabra: Psychology*, 5(1), 50. DOI: <http://doi.org/10.1525/collabra.282>
- Ritov, I., Gati, I., & Tversky, A. (1990). Differential weighting of common and distinctive components. *Journal of Experimental Psychology: General*, 119(1), 30.
- Roads, B. D., & Mozer, M. C. (2019). Obtaining psychological embeddings through joint kernel and metric learning. *Behavior research methods*, 51(5), 2180-2193.
- Rohde, D. L., Gonnerman, L. M., & Plaut, D. C. (2006). An improved model of semantic similarity based on lexical co-occurrence. *Communications of the ACM*, 8(627-633), 116.
- Romney, A. K., Brewer, D. D., & Batchelder, W. H. (1993). Predicting clustering from semantic structure. *Psychological Science*, 4, 28–34. <http://dx.doi.org/10.1111/j.1467-9280.1993.tb00552.x>
- Schlechtweg, D., Oguz, C., & Shculte im Walde, S. (2019). Second-order Co-occurrence Sensitivity of Skip-Gram with Negative Sampling. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (pp. 24-30). Association for Computational Linguistics.
- Schnabel, T., Labutov, I., Mimno, D., & Joachims, T. (2015). Evaluation methods for unsupervised word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* (pp. 298-307).
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. II. *Psychometrika*, 27(3), 219-246.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science.

Science, 237(4820), 1317-1323.

Shepard, R. (2004). How a cognitive psychologist came to seek universal laws. *Psychonomic*

Bulletin & Review, 11(1), 1–23.

Shepard, R. N., & Arabie, P. (1979). Additive clustering representations of similarities as

combinations of discrete overlapping properties. *Psychological Review*, 86, 87-123.

Spearman, C. (1910). Correlation calculated with faulty data. *British Journal of Psychology*, 3,

271-295.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference.

Behavioral and brain sciences, 24(4), 629.

Tissier, J., Gravier, C., & Habrard, A. (2019). Near-lossless binarization of word embeddings. In

Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 33, pp. 7104-7111).

Tsvetkov, Y., Faruqui, M., Ling, W., Lample, G., & Dyer, C. (2015). Evaluation of word vector

representations by subspace alignment. In *Proceedings of the 2015 Conference on*

Empirical Methods in Natural Language Processing (pp. 2049-2054).

Tversky, A. (1969). Intransitivity of preferences. *Psychological review*, 76(1), 31.

Tversky, A. (1977). Features of similarity. *Psychological review*, 84(4), 327.

Tversky, A., & Kahneman, D. (1983). Extensional Versus Intuitive Reasoning: The Conjunction

Fallacy in Probability Judgment. *Psychological Review*, 90(4), 293–315.

<https://doi.org/10.1007/BF00712851>

- Vankrunkelsven, H., Verheyen, S., Storms, G., & De Deyne, S. (2018). Predicting lexical norms: A comparison between a word association model and text-based word co-occurrence models. *Journal of cognition*, 1(1), 45. doi: <http://doi.org/10.5334/joc.50>
- Verheyen, S., Voorspoels, W., Vanpaemel, W., & Storms, G. (2016). Caveats for the spatial arrangement method: Comment on Hout, Goldinger, and Ferguson (2013). *Journal of Experimental Psychology: General*, 145, 376–382. <http://dx.doi.org/10.1037/a0039758>
- Vulić, I., Baker, S., Ponti, E. M., Petti, U., Leviant, I., Wing, K., ... & Reichart, R. (2020). Multi-SimLex: A Large-Scale Evaluation of Multilingual and Cross-Lingual Lexical Semantic Similarity. *arXiv preprint arXiv:2003.04866*.
- Wieting, J., Bansal, M., Gimpel, K., & Livescu, K. (2015). From paraphrase database to compositional paraphrase model and back. *Transactions of the Association for Computational Linguistics*, 3, 345-358.
- Wnuk, E., & Majid, A. (2014). Revisiting the limits of language: The odor lexicon of Maniq. *Cognition*, 131(1), 125-138.

Appendix A

Furniture	Clothing	Birds	Vegetables	Sports	Vehicles	Fruit	Professions
bed	beanie	blackbird	artichoke	archery	airplane	apple	accountant
bench	belt	canary	asparagus	badminton	bicycle	apricot	actor
bookcase	blouse	chicken	beans	ballet	boat	banana	architect
cabinet	boots	crow	beet	baseball	bus	blueberry	baker
chair	boxers	cuckoo	broccoli	basketball	car	cherry	butcher
clock	bra	dove	cabbage	billiards	carriage	coconut	carpenter
cushion	coat	duck	carrot	boxing	cart	date	cashier
desk	dress	eagle	cauliflower	chess	helicopter	grapefruit	clerk
dresser	gloves	falcon	celery	cycling	hovercraft	grapes	cook
footstool	gown	heron	corn	fencing	jeep	honeydew	dentist
lamp	hat	magpie	cucumber	fishing	motorcycle	lemon	doctor
mirror	jacket	ostrich	eggplant	golfing	rocket	melon	educator
piano	jeans	owl	lettuce	gymnastics	scooter	olive	fireman
radio	mittens	parakeet	onion	handball	skateboard	peach	judge
rug	overalls	parrot	potato	hockey	sled	pear	lawyer
sofa	pajamas	peacock	radish	judo	submarine	pineapple	manager
stove	panties	pelican	spinach	rowing	tractor	plum	minister
table	pants	penguin	squash	rugby	trailer	pomegranate	pharmacist
telephone	scarf	pheasant	tomato	running	train	strawberry	physiotherapist
vase	shorts	pigeon	turnip	sailing	trolley	tangerine	pilot
	skirt	robin		skiing	truck	watermelon	plumber
	sneakers	rooster		soccer	van		policeman
	socks	seagull		squash			postman
	suit	sparrow		surfing			psychologist
	sweater	stork		swimming			secretary
	sweatshirt	swallow		tennis			stewardess
	tie	swan		volleyball			teacher
	tracksuit	turkey		walking			veterinarian
	tuxedo	vulture					
		woodpecker					