

Problems With Evaluation of Word Embeddings Using Word Similarity Tasks

Manaal Faruqui¹ Yulia Tsvetkov¹ Pushpendre Rastogi² Chris Dyer¹

¹Language Technologies Institute, Carnegie Mellon University

²Department of Computer Science, Johns Hopkins University

{mfaruqui, ytsvetko, cdyer}@cs.cmu.edu, pushpendre@jhu.edu

Abstract

Lacking standardized extrinsic evaluation methods for vector representations of words, the NLP community has relied heavily on *word similarity* tasks as a proxy for intrinsic evaluation of word vectors. Word similarity evaluation, which correlates the distance between vectors and human judgments of “semantic similarity” is attractive, because it is computationally inexpensive and fast. In this paper we present several problems associated with the evaluation of word vectors on word similarity datasets, and summarize existing solutions. Our study suggests that the use of word similarity tasks for evaluation of word vectors is not sustainable and calls for further research on evaluation methods.

1 Introduction

Despite the ubiquity of word vector representations in NLP, there is no consensus in the community on what is the best way for evaluating word vectors. The most popular intrinsic evaluation task is the *word similarity* evaluation. In word similarity evaluation, a list of pairs of words along with their similarity rating (as judged by human annotators) is provided. The task is to measure how well the notion of word similarity according to humans is captured by the word vector representations. Table 1 shows some word pairs along with their similarity judgments from WS-353 (Finkelstein et al., 2002), a popular word similarity dataset.

Let a, b be two words, and $\mathbf{a}, \mathbf{b} \in \mathbb{R}^D$ be their corresponding word vectors in a D -dimensional vector space. Word similarity in the vector-space can be obtained by computing the cosine similar-

Word ₁	Word ₂	Similarity score [0,10]
love	sex	6.77
stock	jaguar	0.92
money	cash	9.15
development	issue	3.97
lad	brother	4.46

Table 1: Sample word pairs along with their human similarity judgment from WS-353.

ity between the word vectors of a pair of words:

$$\text{cosine}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \quad (1)$$

where, $\|\mathbf{a}\|$ is the ℓ_2 -norm of the vector, and $\mathbf{a} \cdot \mathbf{b}$ is the dot product of the two vectors. Once the vector-space similarity between the words is computed, we obtain the lists of pairs of words sorted according to vector-space similarity, and human similarity. Computing Spearman’s correlation (Myers and Well, 1995) between these ranked lists provides some insight into how well the learned word vectors capture intuitive notions of word similarity.

Word similarity evaluation is attractive, because it is computationally inexpensive and fast, leading to faster prototyping and development of word vector models. The origin of word similarity tasks can be tracked back to Rubenstein and Goode-nough (1965) who constructed a list of 65 word pairs with annotations of human similarity judgment. They created this dataset to validate the veracity of the distributional hypothesis (Harris, 1954) according to which the meaning of words is evidenced by the context they occur in. They found a positive correlation between contextual similarity and human-annotated similarity of word pairs. Since then, the lack of a standard evaluation method for word vectors has led to the creation of several *ad hoc* word similarity datasets. Table 2 provides a list of such benchmarks obtained from `wordvectors.org` (Faruqui and Dyer, 2014a).

Dataset	Word pairs	Reference
RG	65	Rubenstein and Goodenough (1965)
MC	30	Miller and Charles (1991)
WS-353	353	Finkelstein et al. (2002)
YP-130	130	Yang and Powers (2006)
MTurk-287	287	Radinsky et al. (2011)
MTurk-771	771	Halawi et al. (2012)
MEN	3000	Bruni et al. (2012)
RW	2034	Luong et al. (2013)
Verb	144	Baker et al. (2014)
SimLex	999	Hill et al. (2014)

Table 2: Word similarity datasets.

In this paper, we give a comprehensive analysis of the problems that are associated with the evaluation of word vector representations using word similarity tasks.¹ We survey existing literature to construct a list of such problems and also summarize existing solutions to some of the problems. Our findings suggest that word similarity tasks are not appropriate for evaluating word vector representations, and call for further research on better evaluation methods

2 Problems

We now discuss the major issues with evaluation of word vectors using word similarity tasks, and present existing solutions (if available) to address them.

2.1 Subjectivity of the task

The notion of word similarity is subjective and is often confused with relatedness. For example, *cup*, and *coffee* are related to each other, but not similar. *Coffee* refers to a plant (a living organism) or a hot brown drink, whereas *cup* is a man-made object, which contains liquids, often coffee. Nevertheless, *cup* and *coffee* are rated more similar than pairs such as *car* and *train* in WS-353 (Finkelstein et al., 2002). Such anomalies are also found in recently constructed datasets like MEN (Bruni et al., 2012). Thus, such datasets unfairly penalize word vector models that capture the fact that *cup* and *coffee* are dissimilar.

¹An alternative to correlation-based word similarity evaluation is the *word analogy* task, where the task is to find the missing word b^* in the relation: a is to a^* as b is to b^* , where a, a^* are related by the same relation as a, a^* . For example, $king : man :: queen : woman$. Mikolov et al. (2013b) showed that this problem can be solved using the vector offset method: $\mathbf{b}^* \approx \mathbf{b} - \mathbf{a} + \mathbf{a}^*$. Levy and Goldberg (2014a) show that solving this equation is equivalent to computing a linear combination of word similarities between the query word b^* , with the given words a, b , and a^* . Thus, the results we present in this paper naturally extend to the word analogy tasks.

In an attempt to address this limitation, Agirre et al. (2009) divided WS-353 into two sets containing word pairs exhibiting only either similarity or relatedness. Recently, Hill et al. (2014) constructed a new word similarity dataset (SimLex), which captures the degree of similarity between words, and related words are considered dissimilar. Even though it is useful to separate the concept of similarity and relatedness, it is not clear as to which one should the word vector models be expected to capture.

2.2 Semantic or task-specific similarity?

Distributional word vector models capture some aspect of word co-occurrence statistics of the words in a language (Levy and Goldberg, 2014b; Levy et al., 2015). Therefore, to the extent these models produce semantically coherent representations, it can be seen as evidence of the distributional hypothesis of Harris (1954). Thus, word embeddings like Skip-gram, CBOW, Glove, LSA (Turney and Pantel, 2010; Mikolov et al., 2013a; Pennington et al., 2014) which are trained on word co-occurrence counts can be expected to capture semantic word similarity, and hence can be evaluated on word similarity tasks.

Word vector representations which are trained as part of a neural network to solve a particular task (apart from word co-occurrence prediction) are called distributed word embeddings (Collobert and Weston, 2008), and they are task-specific in nature. These embeddings capture task-specific word similarity, for example, if the task is of POS tagging, two nouns *cat* and *man* might be considered similar by the model, even though they are not semantically similar. Thus, evaluating such task-specific word embeddings on word similarity can unfairly penalize them. This raises the question: what kind of word similarity should be captured by the model?

2.3 No standardized splits & overfitting

To obtain generalizable machine learning models, it is necessary to make sure that they do not overfit to a given dataset. Thus, the datasets are usually partitioned into a training, development and test set on which the model is trained, tuned and finally evaluated, respectively (Manning and Schütze, 1999). Existing word similarity datasets are not partitioned into training, development and test sets. Therefore, optimizing the word vectors to perform better at a word similarity task implic-

itly *tunes on the test set* and overfits the vectors to the task. On the other hand, if researchers decide to perform their own splits of the data, the results obtained across different studies can be incomparable. Furthermore, the average number of word pairs in the word similarity datasets is small (≈ 781 , cf. Table 2), and partitioning them further into smaller subsets may produce unstable results.

We now present some of the solutions suggested by previous work to avoid overfitting of word vectors to word similarity tasks. Faruqui and Dyer (2014b), and Lu et al. (2015) evaluate the word embeddings exclusively on word similarity and word analogy tasks. Faruqui and Dyer (2014b) tune their embedding on one word similarity task and evaluate them on all other tasks. This ensures that their vectors are being evaluated on held-out datasets. Lu et al. (2015) propose to directly evaluate the generalization of a model by measuring the performance of a single model on a large gamut of tasks. This evaluation can be performed in two different ways: (1) choose the hyperparameters with best average performance across all tasks, (2) choose the hyperparameters that beat the baseline vectors on most tasks.² By selecting the hyperparameters that perform well across a range of tasks, these methods ensure that the obtained vectors are generalizable. Stratos et al. (2015) divided each word similarity dataset individually into tuning and test set and reported results on the test set.

2.4 Low correlation with extrinsic evaluation

Word similarity evaluation measures how well the notion of word similarity according to humans is captured in the vector-space word representations. Word vectors that can capture word similarity might be expected to perform well on tasks that require a notion of explicit semantic similarity between words like paraphrasing, entailment. However, it has been shown that no strong correlation is found between the performance of word vectors on word similarity and extrinsic evaluation NLP tasks like text classification, parsing, sentiment analysis (Tsvetkov et al., 2015; Schnabel et al., 2015).³ An absence of strong correlation between the word similarity evaluation and downstream tasks calls for alternative approaches

²Baseline vectors can be any off-the-shelf vector models.

³In these studies, extrinsic evaluation tasks are those tasks that use the dimensions of word vectors as features in a machine learning model. The model learns weights for how important these features are for the extrinsic task.

to evaluation.

2.5 Absence of statistical significance

There has been a consistent omission of statistical significance for measuring the difference in performance of two vector models on word similarity tasks. Statistical significance testing is important for validating metric gains in NLP (Berg-Kirkpatrick et al., 2012; Søgaard et al., 2014), specifically while solving non-convex objectives where results obtained due to optimizer instability can often lead to incorrect inferences (Clark et al., 2011). The problem of statistical significance in word similarity evaluation was first systematically addressed by Shalaby and Zadrozny (2015), who used Steiger’s test (Steiger, 1980)⁴ to compute how significant the difference between rankings produced by two different models is against the gold ranking. However, their method needs explicit ranked list of words produced by the models and cannot work when provided only with the correlation ratio of each model with the gold ranking. This problem was solved by Rastogi et al. (2015), which we describe next.

Rastogi et al. (2015) observed that the improvements shown on small word similarity task datasets by previous work were insignificant. We now briefly describe the method presented by them to compute statistical significance for word similarity evaluation. Let A and B be the rankings produced by two word vector models over a list of words pairs, and T be the human annotated ranking. Let r_{AT} , r_{BT} and r_{AB} denote the Spearman’s correlation between $A : T$, $B : T$ and $A : B$ resp. and \hat{r}_{AT} , \hat{r}_{BT} and \hat{r}_{AB} be their empirical estimates. Rastogi et al. (2015) introduce $\sigma_{p_0}^r$ as the minimum required difference for significance (MRDS) which satisfies the following:

$$(r_{AB} < r) \wedge (|\hat{r}_{BT} - \hat{r}_{AT}| < \sigma_{p_0}^r) \implies pval > p_0 \quad (2)$$

Here $pval$ is the probability of the test statistic under the null hypothesis that $r_{AT} = r_{BT}$ found using the Steiger’s test. The above conditional ensures that if the empirical difference between the rank correlations of the scores of the competing methods to the gold ratings is less than $\sigma_{p_0}^r$ then either the true correlation between the competing methods is greater than r , or the null hypothesis of no difference has p-value greater than p_0 . $\sigma_{p_0}^r$

⁴A quick tutorial on Steiger’s test & scripts: <http://www.philippsinger.info/?p=347>

depends on the size of the dataset, p_0 and r and Rastogi et al. (2015) present its values for common word similarity datasets. Reporting statistical significance in this way would help estimate the differences between word vector models.

2.6 Frequency effects in cosine similarity

The most common method of measuring the similarity between two words in the vector-space is to compute the cosine similarity between the corresponding word vectors. Cosine similarity implicitly measures the similarity between two unit-length vectors (eq. 1). This prevents any biases in favor of frequent words which are longer as they are updated more often during training (Turian et al., 2010).

Ideally, if the geometry of embedding space is primarily driven by semantics, the relatively small number of frequent words should be evenly distributed through the space, while large number of rare words should cluster around related, but more frequent words. However, it has been shown that vector-spaces contain *hubs*, which are vectors that are close to a large number of other vectors in the space (Radovanović et al., 2010). This problem manifests in word vector-spaces in the form of words that have high cosine similarity with a large number of other words (Dinu et al., 2014). Schnabel et al. (2015) further refine this *hubness* problem to show that there exists a power-law relationship between the frequency-rank⁵ of a word and the frequency-rank of its neighbors. Specifically, they showed that the average rank of the 1000 nearest neighbors of a word follows:

$$\text{nn-rank} \approx 1000 \cdot \text{word-rank}^{0.17} \quad (3)$$

This shows that pairs of words which have similar frequency will be closer in the vector-space, thus showing higher word similarity than they should according to their word meaning. Even though newer datasets of word similarity sample words from different frequency bins (Luong et al., 2013; Hill et al., 2014), this still does not solve the problem that cosine similarity in the vector-space gets polluted by frequency-based effects. Different distance normalization schemes have been proposed to downplay the frequency/hubness effect when computing nearest neighbors in the vector space (Dinu et al., 2014; Tomašev et al., 2011),

⁵The rank of a word in vocabulary of the corpus sorted in decreasing order of frequency.

but their applicability as an absolute measure of distance for word similarity tasks still needs to be investigated.

2.7 Inability to account for polysemy

Many words have more than one meaning in a language. For example, the word *bank* can either correspond to a financial institution or to the land near a river. However in WS-353, *bank* is given a similarity score of 8.5/10 to *money*, signifying that *bank* is a financial institution. Such an assumption of one sense per word is prevalent in many of the existing word similarity tasks, and it can incorrectly penalize a word vector model for capturing a specific sense of the word absent in the word similarity task.

To account for sense-specific word similarity, Huang et al. (2012) introduced the Stanford contextual word similarity dataset (SCWS), in which the task is to compute similarity between two words given the contexts they occur in. For example, the words *bank* and *money* should have a low similarity score given the contexts: “*along the east bank of the river*”, and “*the basis of all money laundering*”. Using cues from the word’s context, the correct word-sense can be identified and the appropriate word vector can be used. Unfortunately, word senses are also ignored by majority of the frequently used word vector models like Skip-gram and Glove. However, there has been progress on obtaining multiple vectors per word-type to account for different word-senses (Reisinger and Mooney, 2010; Huang et al., 2012; Neelakantan et al., 2014; Jauhar et al., 2015; Rothe and Schütze, 2015).

3 Conclusion

In this paper we have identified problems associated with word similarity evaluation of word vector models, and reviewed existing solutions wherever possible. Our study suggests that the use of word similarity tasks for evaluation of word vectors can lead to incorrect inferences and calls for further research on evaluation methods.

Until a better solution is found for intrinsic evaluation of word vectors, we suggest task-specific evaluation: word vector models should be compared on how well they can perform on a downstream NLP task. Although task-specific evaluation produces different rankings of word vector models for different tasks (Schnabel et al., 2015),

this is not necessarily a problem because different vector models capture different types of information which can be more or less useful for a particular task.

References

- Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Pasca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and wordnet-based approaches. In *Proc. of NAACL*.
- Simon Baker, Roi Reichart, and Anna Korhonen. 2014. An unsupervised model for instance level subcategory acquisition. In *Proc. of EMNLP*.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. An empirical investigation of statistical significance in nlp. In *Proc. of EMNLP*.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proc. of ACL*.
- Jonathan H Clark, Chris Dyer, Alon Lavie, and Noah A Smith. 2011. Better hypothesis testing for statistical machine translation: Controlling for optimizer instability. In *Proc. of ACL*.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In *Proc. of ICML*.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. 2014. Improving zero-shot learning by mitigating the hubness problem. *arXiv preprint arXiv:1412.6568*.
- Manaal Faruqui and Chris Dyer. 2014a. Community evaluation and exchange of word vectors at word-vectors.org. In *Proc. of ACL: System Demo*.
- Manaal Faruqui and Chris Dyer. 2014b. Improving vector space word representations using multilingual correlation. In *Proc. of EACL*.
- Lev Finkelstein, Evgeniy Gabilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2002. Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20(1).
- Guy Halawi, Gideon Dror, Evgeniy Gabilovich, and Yehuda Koren. 2012. Large-scale learning of word relatedness with constraints. In *Proc. of SIGKDD*.
- Zellig Harris. 1954. Distributional structure. *Word*, 10(23):146–162.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2014. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *CoRR*, abs/1408.3456.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proc. of ACL*.
- Sujay Kumar Jauhar, Chris Dyer, and Eduard Hovy. 2015. Ontologically grounded multi-sense representation learning for semantic vector space models. In *Proc. NAACL*.
- Omer Levy and Yoav Goldberg. 2014a. Linguistic regularities in sparse and explicit word representations. In *Proc. of CoNLL*.
- Omer Levy and Yoav Goldberg. 2014b. Neural word embedding as implicit matrix factorization. In *Proc. of NIPS*.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *TACL*, 3:211–225.
- Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Deep multilingual correlation for improved word embeddings. In *Proc. of NAACL*.
- Minh-Thang Luong, Richard Socher, and Christopher D. Manning. 2013. Better word representations with recursive neural networks for morphology. In *Proc. of CoNLL*.
- Christopher D. Manning and Hinrich Schütze. 1999. *Foundations of statistical natural language processing*. MIT Press, Cambridge, MA, USA.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. In *Proc. of NAACL*.
- George Miller and Walter Charles. 1991. Contextual correlates of semantic similarity. In *Language and Cognitive Processes*, pages 1–28.
- Jerome L. Myers and Arnold D. Well. 1995. *Research Design & Statistical Analysis*. Routledge.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proc. of EMNLP*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proc. of EMNLP*.
- Kira Radinsky, Eugene Agichtein, Evgeniy Gabilovich, and Shaul Markovitch. 2011. A word at a time: computing word relatedness using temporal semantic analysis. In *Proc. of WWW*.

- Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. Hubs in space: Popular nearest neighbors in high-dimensional data. *The Journal of Machine Learning Research*, 11:2487–2531.
- Pushpendre Rastogi, Benjamin Van Durme, and Raman Arora. 2015. Multiview lsa: Representation learning via generalized cca. In *Proc. of NAACL*.
- Joseph Reisinger and Raymond J. Mooney. 2010. Multi-prototype vector-space models of word meaning. In *Proc. of NAACL*.
- Sascha Rothe and Hinrich Schütze. 2015. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. In *Proc. of ACL*.
- Herbert Rubenstein and John B. Goodenough. 1965. Contextual correlates of synonymy. *Commun. ACM*, 8(10).
- Tobias Schnabel, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proc. of EMNLP*.
- Walid Shalaby and Wlodek Zadrozny. 2015. Measuring semantic relatedness using mined semantic analysis. *CoRR*, abs/1512.03465.
- Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Héctor Martínez Alonso. 2014. What’s in a p-value in nlp? In *Proc. of CoNLL*.
- James H Steiger. 1980. Tests for comparing elements of a correlation matrix. *Psychological bulletin*, 87(2):245.
- Karl Stratos, Michael Collins, and Daniel Hsu. 2015. Model-based word embeddings from decompositions of count matrices. In *Proc. of ACL*.
- Nenad Tomašev, Miloš Radovanovic, Dunja Mladenic, and Mirjana Ivanovic. 2011. A probabilistic approach to nearest-neighbor classification: Naive hubness bayesian knn. In *Proc. of CIKM*.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *Proc. of EMNLP*.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proc. of ACL*.
- Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning : Vector space models of semantics. *JAIR*, pages 141–188.
- Dongqiang Yang and David M. W. Powers. 2006. Verb similarity on the taxonomy of wordnet. In *3rd International WordNet Conference*.