

Query Expansion with Locally-Trained Word Embeddings

Fernando Diaz

Microsoft

fdiaz@microsoft.com

Bhaskar Mitra

Microsoft

bmitra@microsoft.com

Nick Craswell

Microsoft

nickr@microsoft.com

Abstract

Continuous space word embeddings have received a great deal of attention in the natural language processing and machine learning communities for their ability to model term similarity and other relationships. We study the use of term relatedness in the context of query expansion for *ad hoc* information retrieval. We demonstrate that word embeddings such as word2vec and GloVe, when trained globally, underperform corpus and query specific embeddings for retrieval tasks. These results suggest that other tasks benefiting from global embeddings may also benefit from local embeddings.

1 Introduction

Continuous space embeddings such as word2vec (Mikolov et al., 2013b) or GloVe (Pennington et al., 2014a) project terms in a vocabulary to a dense, lower dimensional space. Recent results in the natural language processing community demonstrate the effectiveness of these methods for analogy and word similarity tasks. In general, these approaches provide global representations of words; each word has a fixed representation, regardless of any discourse context. While a global representation provides some advantages, language use can vary dramatically by topic. For example, ambiguous terms can easily be disambiguated given local information in immediately surrounding words (Harris, 1954; Yarowsky, 1993). The window-based training of word2vec style algorithms exploits this distributional property.

A global word embedding, even when

trained using local windows, risks capturing only coarse representations of those topics dominant in the corpus. While a particular embedding may be appropriate for a specific word within a sentence-length context globally, it may be entirely inappropriate within a specific topic. Gale *et al.* refer to this as the ‘one sense per discourse’ property (Gale et al., 1992). Previous work by Yarowsky demonstrates that this property can be successfully combined with information from nearby terms for word sense disambiguation (Yarowsky, 1995). Our work extends this approach to word2vec-style training in the context word similarity.

For many tasks that require topic-specific linguistic analysis, we argue that topic-specific representations should outperform global representations. Indeed, it is difficult to imagine a natural language processing task that would *not* benefit from an understanding of the local topical structure. Our work focuses on a query expansion, an information retrieval task where we can study different lexical similarity methods with an extrinsic evaluation metric (i.e. retrieval metrics). Recent work has demonstrated that similarity based on global word embeddings can be used to outperform classic pseudo-relevance feedback techniques (Sordani et al., 2014; al Masri et al., 2016).

We propose that embeddings be learned on topically-constrained corpora, instead of large topically-unconstrained corpora. In a retrieval scenario, this amounts to retraining an embedding on documents related to the topic of the query. We present local embeddings which capture the nuances of topic-specific language better than global embeddings. There is substantial evidence that global methods underperform local methods for information re-

trieval tasks such as query expansion (Xu and Croft, 1996), latent semantic analysis (Hull, 1994; Schütze et al., 1995; Singhal et al., 1997), cluster-based retrieval (Tombros and van Rijsbergen, 2001; Tombros et al., 2002; Willett, 1985), and term clustering (Attar and Fraenkel, 1977). We demonstrate that the same holds true when using word embeddings for text retrieval.

2 Motivation

For the purpose of motivating our approach, we will restrict ourselves to word2vec although other methods behave similarly (Levy and Goldberg, 2014). These algorithms involve discriminatively training a neural network to predict a word given small set of context words. More formally, given a target word w and observed context c , the instance loss is defined as,

$$\begin{aligned} \ell(w, c) = & \log \sigma(\phi(w) \cdot \psi(c)) \\ & + \eta \cdot \mathbb{E}_{\bar{w} \sim \theta_C} [\log \sigma(-\phi(w) \cdot \psi(\bar{w}))] \end{aligned}$$

where $\phi : \mathcal{V} \rightarrow \mathbb{R}^k$ projects a term into a k -dimensional embedding space, $\psi : \mathcal{V}^m \rightarrow \mathbb{R}^k$ projects a set of m terms into a k -dimensional embedding space, and \bar{w} is a randomly sampled ‘negative’ context. The parameter η controls the sampling of random negative terms. These matrices are estimated over a set of contexts sampled from a large corpus and minimize the expected loss,

$$\mathcal{L}_c = \mathbb{E}_{w, c \sim p_c} [\ell(w, c)] \quad (1)$$

where p_c is the distribution of word-context pairs in the training corpus and can be estimated from corpus statistics.

While using corpus statistics may make sense absent any other information, oftentimes we know that our analysis will be topically constrained. For example, we might be analyzing the ‘sports’ documents in a collection. The language in this domain is more specialized and the distribution over word-context pairs is unlikely to be similar to $p_c(w, c)$. In fact, prior work in information retrieval suggests that documents on subtopics in a collection have very different unigram distributions compared to the whole corpus (Cronen-Townsend et al., 2002). Let $p_t(w, c)$ be the probability




Figure 1: Importance weights for terms occurring in documents related to ‘argentina pegging dollar’ relative to frequency in gigaword.

of observing a word-context pair conditioned on the topic t . The expected loss under this distribution is (Shimodaira, 2000),

$$\mathcal{L}_t = \mathbb{E}_{w, c \sim p_c} \left[\frac{p_t(w, c)}{p_c(w, c)} \ell(w, c) \right] \quad (2)$$

In general, if our corpus consists of sufficiently diverse data (e.g. Wikipedia), the support of $p_t(w, c)$ is much smaller than and contained in that of $p_c(w, c)$. The loss, ℓ , of a context that occurs more frequently in the topic, will be amplified by the importance weight $\omega = \frac{p_t(w, c)}{p_c(w, c)}$. Because topics require specialized language, this is likely to occur; at the same time, these contexts are likely to be underemphasized in training a model according to Equation 1.

In order to quantify this, we took a topic from a TREC *ad hoc* retrieval collection (see Section 5 for details) and computed the importance weight for each term occurring in the set of on-topic documents. The histogram of weights ω is presented in Figure 1. While larger probabilities are expected since the size of a topic-constrained vocabulary is smaller, there are a non-trivial number of terms with much larger importance weights. If the loss, $\ell(w)$, of a word2vec embedding is worse for these words with low $p_c(w)$, then we expect these errors to be exacerbated for the topic.

Of course, these highly weighted terms may have a low value for $p_t(w)$ but a very high value relative to the corpus. We can adjust the




Figure 2: Pointwise Kullback-Leibler divergence for terms occurring in documents related to ‘argentina pegging dollar’ relative to frequency in gigaword.

weights by considering the pointwise Kullback-Leibler divergence for each word w ,

$$D_w(p_t||p_c) = p_t(w) \log \frac{p_t(w)}{p_c(w)} \quad (3)$$

Words which have a much higher value of $p_t(w)$ than $p_c(w)$ and have a high absolute value of $p_t(w)$ will have high pointwise KL divergence. Figure 2 shows the divergences for the top 100 most frequent terms in $p_t(w)$. The higher ranked terms (i.e. good query expansion candidates) tend to have much higher probabilities than found in $p_c(w)$. If the loss on those words is large, this may result in poor embeddings for the most important words for the topic.

A dramatic change in distribution between the corpus and the topic has implications for performance precisely because of the objective used by word2vec (i.e. Equation 1). The training emphasizes word-context pairs occurring with high frequency in the corpus. We will demonstrate that, even with heuristic down-sampling of frequent terms in word2vec, these techniques result in inferior performance for specific topics.

Thus far, we have sketched out why using the corpus distribution for a specific topic may result in undesirable outcomes. However, it is even unclear that $p_t(w|c) = p_c(w|c)$. In fact, we suspect that $p_t(w|c) \neq p_c(w|c)$ because of the ‘one sense per discourse’ claim (Gale et al., 1992). We can qualitatively observe the difference in $p_c(w|c)$ and $p_t(w|c)$ by training

global	local
cutting	tax
squeeze	deficit
reduce	vote
slash	budget
reduction	reduction
spend	house
lower	bill
halve	plan
soften	spend
freeze	billion

Figure 3: Terms similar to ‘cut’ for a word2vec model trained on a general news corpus and another trained only on documents related to ‘gasoline tax’.

two word2vec models: the first on the large, generic Gigaword corpus and the second on a topically-constrained subset of the gigaword. We present the most similar terms to ‘cut’ using both a global embedding and a topic-specific embedding in Figure 3. In this case, the topic is ‘gasoline tax’. As we can see, the ‘tax cut’ sense of ‘cut’ is emphasized in the topic-specific embedding.

3 Local Word Embeddings

The previous section described several reasons why a global embedding may result in over-general word embeddings. In order to perform topic-specific training, we need a set of topic-specific documents. In information retrieval scenarios users rarely provide the system with examples of topic-specific documents, instead providing a small set of keywords.

Fortunately, we can use information retrieval techniques to generate a query-specific set of topical documents. Specifically, we adopt a language modeling approach to do so (Croft and Lafferty, 2003). In this retrieval model, each document is represented as a maximum likelihood language model estimated from document term frequencies. Query language models are estimated similarly, using term frequency in the query. A document score then, is the Kullback-Leibler divergence between the query and document language

models,

$$D(p_q \| p_d) = \sum_{w \in \mathcal{V}} p_q(w) \log \frac{p_q(w)}{p_d(w)} \quad (4)$$

Documents whose language models are more similar to the query language model will have a *lower* KL divergence score. For consistency with prior work, we will refer to this as the query likelihood score of a document.

The scores in Equation 4 can be passed through a softmax function to derive a multinomial over the entire corpus (Lavrenko and Croft, 2001),

$$p(d) = \frac{\exp(-D(p_q \| p_d))}{\sum_{d'} \exp(-D(p_q \| p_{d'}))} \quad (5)$$

Recall in Section 2 that training a word2vec model weights word-context pairs according to the corpus frequency. Our query-based multinomial, $p(d)$, provides a weighting function capturing the documents relevant to this topic. Although an estimation of the topic-specific documents from a query will be imprecise (i.e. some nonrelevant documents will be scored highly), the language use tends to be consistent with that found in the known relevant documents.

We can train a local word embedding using an arbitrary optimization method by sampling documents from $p(d)$ instead of uniformly from the corpus. In this work, we use word2vec, although any method that operates on a sample of documents can be used.

4 Query Expansion with Word Embeddings

When using language models for retrieval, query expansion involves estimating an alternative to p_q . Specifically, when each expansion term is associated with a weight, we normalize these weights to derive the expansion language model, p_{q+} . This language model is then interpolated with the original query model,

$$p_q^1(w) = \lambda p_q(w) + (1 - \lambda) p_{q+}(w) \quad (6)$$

This interpolated language model can then be used with Equation 4 to rank documents (Abdul-Jaleel et al., 2004). We will refer to this as the expanded query score of a document.

Now we turn to using word embeddings for query expansion. Let \mathbf{U} be an $|\mathcal{V}| \times k$ term embedding matrix. If \mathbf{q} is a $|\mathcal{V}| \times 1$ column term vector for a query, then the expansion term weights are $\mathbf{U}\mathbf{U}^T\mathbf{q}$. We then take the top k terms, normalize their weights, and compute $p_{q+}(w)$.

We consider the following alternatives for \mathbf{U} . The first approach is to use a global model trained by sampling documents uniformly. The second approach, which we propose in this paper, is to use a local model trained by sampling documents from $p(d)$.

5 Methods

5.1 Data

To evaluate the different retrieval strategies described in Section 3, we use the following datasets. Two newswire datasets, trec12 and robust, consist of the newswire documents and associated queries from TREC *ad hoc* retrieval evaluations. The trec12 corpus consists of Tipster disks 1 and 2; and the robust corpus consists of Tipster disks 4 and 5. Our third dataset, web, consists of the ClueWeb 2009 Category B Web corpus. For the Web corpus, we only retain documents with a Waterloo spam rank above 70.¹ We present corpus statistics in Table 1.

We consider several publicly available global embeddings. We use four GloVe embeddings of different dimensionality trained on the union of Wikipedia and Gigaword documents.² We use one publicly available word2vec embedding trained on Google News documents.³ We also trained a global embedding for trec12 and robust using the entire corpus. Instead of training a global embedding on the large web collection, we use a GloVe embedding trained on Common Crawl data.⁴

We train local embeddings with word2vec using one of three retrieval sources. First, we consider documents retrieved from the target corpus of the query (i.e. trec12, robust, or web). We also consider training a local embed-

¹<https://plg.uwaterloo.ca/~gvcormac/clueweb09spam/>

²<http://nlp.stanford.edu/data/glove.6B.zip>

³<https://code.google.com/archive/p/word2vec/>

⁴<http://nlp.stanford.edu/data/glove.840B.300d.zip>

	docs	words	queries
trec12	469,949	438,338	150
robust	528,155	665,128	250
web	50,220,423	90,411,624	200
news	9,875,524	2,645,367	-
wiki	3,225,743	4,726,862	-

Table 1: Corpora used for retrieval and local embedding training.

ding by performing a retrieval on large auxiliary corpora. We use the Gigaword corpus as a large auxiliary news corpus. We hypothesize that retrieving from a larger news corpus will provide substantially more local training data than a target retrieval. We also use a Wikipedia snapshot from December 2014. We hypothesize that retrieving from a large, high fidelity corpus will provide cleaner language than that found in lower fidelity target domains such as the web. Table 1 shows the relative magnitude of these auxiliary corpora compared to the target corpora.

All corpora in Table 1 were stopped using the SMART stopword list⁵ and stemmed using the Krovetz algorithm (Krovetz, 1993). We used the Indri implementation for indexing and retrieval.⁶

5.2 Evaluation

We consider several standard retrieval evaluation metrics, including NDCG@10 and interpolated precision at standard recall points (Järvelin and Kekäläinen, 2002; van Rijsbergen, 1979). NDCG@10 provides insight into performance specifically at higher ranks. An interpolated precision recall graph describes system performance throughout the entire ranked list.

5.3 Training

All retrieval experiments were conducted by performing 10-fold cross-validation across queries. Specifically, we cross-validate the number of expansion terms, $k \in \{5, 10, 25, 50, 100, 250, 500\}$, and interpolation weight, $\lambda \in [0, 1]$. For local word2vec training, we cross-validate the learning rate $\alpha \in \{10^{-1}, 10^{-2}, 10^{-3}\}$.

⁵<http://jmlr.csail.mit.edu/papers/volume5/lewis04a/a11-smart-stop-list/english.stop>

⁶<http://www.lemurproject.org/indri/>

All word2vec training used the publicly available word2vec cbow implementation.⁷ When training the local models, we sampled 1000 documents from $p(d)$ with replacement. To compensate for the much smaller corpus size, we ran word2vec training for 80 iterations. Local word2vec models use a fixed embedding dimension of 400 although other choices did not significantly affect our results. Unless otherwise noted, default parameter settings were used.

In our experiments, expanded queries rescore the top 1000 documents from an initial query likelihood retrieval. Previous results have demonstrated that this approach results in performance nearly identical with an expanded retrieval at a much lower cost (Diaz, 2015). Because publicly available embeddings may have tokenization inconsistent with our target corpora, we restricted the vocabulary of candidate expansion terms to those occurring in the initial retrieval. If a candidate term was not found in the vocabulary of the embedding matrix, we searched for the candidate in a stemmed version of the embedding vocabulary. In the event that the candidate term was still not found after this process, we removed it from consideration.

6 Results

We present results for retrieval experiments in Table 2. We find that embedding-based query expansion outperforms our query likelihood baseline across all conditions. When using the global embedding, the news corpora benefit from the various embeddings in different situations. Interestingly, for trec12, using an embedding trained on the target corpus significantly outperforms all other global embeddings, despite using substantially less data to estimate the model. While this performance may be due to the embedding having a tokenization consistent with the target corpus, it may also come from the fact that the corpus is more representative of the target documents than other embeddings which rely on online news or are mixed with non-news content. To some extent this supports our desire to move training closer to the target distribution.

Across all conditions, local embeddings sig-

⁷<https://code.google.com/p/word2vec/>

Table 2: Retrieval results comparing query expansion based on various global and local embeddings. Bolded numbers indicate the best expansion in that class of embeddings. Wilcoxon signed rank test between bolded numbers indicates statistically significant improvements ($p < 0.05$) for all collections.

	QL	global						local		
		wiki+giga				gnews	target	target	giga	wiki
		50	100	200	300	300	400	400	400	400
trec12	0.514	0.518	0.518	0.530	0.531	0.530	0.545	0.535	0.563*	0.523
robust	0.467	0.470	0.463	0.469	0.468	0.472	0.465	0.475	0.517*	0.476
web	0.216	0.227	0.229	0.230	0.232	0.218	0.216	0.234	0.236	0.258*




Figure 4: Interpolated precision-recall curves for query likelihood, the best global embedding, and the best local embedding from Table 2.

nificantly outperform global embeddings for query expansion. For our two news collections, estimating the local model using a retrieval from the larger Gigaword corpus led to substantial improvements. This effect is almost certainly due to the Gigaword corpus being similar in writing style to the target corpus but, at the same time, providing significantly more relevant content (Diaz and Metzler, 2006). As a result, the local embedding is trained using a larger variety of topical material than if it were to use a retrieval from the smaller target corpus. An embedding trained with a retrieval from Wikipedia tended to perform worse most likely because the language is dissimilar from news content. Our web collection, on the other hand, benefitted more from embeddings trained using retrievals from the general Wikipedia corpus. The Gigaword corpus was less useful here because news-style language is almost certainly not representative of general web documents.

Figure 4 presents interpolated precision-recall curves comparing the baseline, the best

global query expansion method, and the best local query expansion method. Interestingly, although global methods achieve strong performance for NDCG@10, these improvements over the baseline are not reflected in our precision-recall curves. Local methods, on the other hand, almost always strictly dominate both the baseline and global expansion across all recall levels.

The results support the hypothesis that local embeddings provide better similarity measures than global embeddings for query expansion. In order to understand why, we first compare the performance differences between local and global embeddings. Figure 2 suggests that we should adopt a local embedding when the local unigram language model deviates from the corpus language model. To test this, we computed the KL divergence between the local unigram distribution, $\sum_d p(w|d)p(d)$, and the corpus unigram language model (Cronen-Townsend et al., 2002). We hypothesize that, when this value is high, the topic language is different from the corpus language and the

Table 3: Kendall’s τ and Spearman’s ρ between improvement in NDCG@10 and local KL divergence with the corpus language model. The improvement is measured for the best local embedding over the best global embedding.

	τ	ρ
trec12	0.0585	0.0798
robust	0.0545	0.0792
web	0.0204	0.0283

global embedding will be inferior to the local embedding. We tested the rank correlation between this KL divergence and the relative performance of the local embedding with respect to the global embedding. These correlations are presented in Table 3. Unfortunately, we find that the correlation is low, although it is positive across collections.

We can also qualitatively analyze the differences in the behavior of the embeddings. If we have access to the set of documents labeled relevant to a query, then we can compute the frequency of terms in this set and consider those terms with high frequency (after stopping and stemming) to be good query expansion candidates. We can then visualize where these terms lie in the global and local embeddings. In Figure 5, we present a two-dimensional projection (van der Maaten and Hinton, 2008) of terms for the query ‘ocean remote sensing’, with those good candidates highlighted. Our projection includes the top 50 candidates by frequency and a sample of terms occurring in the query likelihood retrieval. We notice that, in the global embedding, the good candidates are spread out amongst poorer candidates. By contrast, the local embedding clusters the candidates in general but also situates them closely around the query. As a result, we suspect that the similar terms extracted from the local embedding are more likely to include these good candidates.

7 Discussion

The success of local embeddings on this task should alarm natural language processing researchers using global embeddings as a representational tool. For one, the approach of learning from vast amounts of data is only ef-

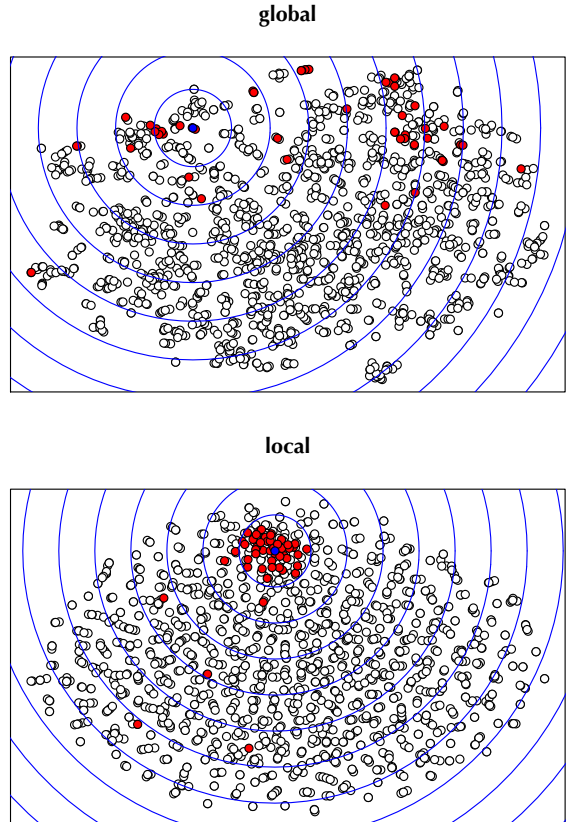


Figure 5: Global versus local embedding of highly relevant terms. Each point represents a candidate expansion term. Red points have high frequency in the relevant set of documents. White points have low or no frequency in the relevant set of documents. The blue point represents the query. Contours indicate distance from the query.

fective if the data is appropriate for the task at hand. And, when provided, much smaller high-quality data can provide much better performance. Beyond this, our results suggest that the approach of estimating global representations, while computationally convenient, may overlook insights possible at query time, or evaluation time in general. A similar local embedding approach can be adopted for any natural language processing task where topical locality is expected and can be estimated. Although we used a query to re-weight the corpus in our experiments, we could just as easily use alternative contextual information (e.g. a sentence, paragraph, or document) in other tasks.

Despite these strong results, we believe that

there are still some open questions in this work. First, although local embeddings provide effectiveness gains, they can be quite inefficient compared to global embeddings. We believe that there is opportunity to improve the efficiency by considering offline computation of local embeddings at a coarser level than queries but more specialized than the corpus. If the retrieval algorithm is able to select the appropriate embedding at query time, we can avoid training the local embedding. Second, although our supporting experiments (Table 3, Figure 5) add some insight into our intuition, the results are not strong enough to provide a solid explanation. Further theoretical and empirical analysis is necessary.

8 Related Work

Topical adaptation of models The shortcomings of learning a single global vector representation, especially for polysemic words, have been pointed out before (Reisinger and Mooney, 2010b). The problem can be addressed by training a global model with multiple vector embeddings per word (Reisinger and Mooney, 2010a; Huang et al., 2012) or topic-specific embeddings (Liu et al., 2015). The number of senses for each word may be fixed (Neelakantan et al., 2015), or determined using class labels (Trask et al., 2015). However, to the best of our knowledge, this is the first time that training topic-specific word embeddings has been explored.

Several methods exist in the language modeling community for topic-dependent adaptation of language models (Bellegarda, 2004). These can lead to performance improvements in tasks such as machine translation (Zhao et al., 2004) and speech recognition (Nanjo and Kawahara, 2004). Topic-specific data may be gathered in advance, by identifying corpus of topic-specific documents. It may also be gathered during the discourse, using multiple hypotheses from N-best lists as a source of topic-specific language. Then a topic-specific language model is trained (or the global model is adapted) online using the topic-specific training data. A topic-dependent model may be combined with the global model using linear interpolation (Iyer and Ostendorf, 1999) or other more sophisticated approaches (Fed-

erico, 1996; Kuhn and De Mori, 1990). Similarly to the adaptation work, we use topic-specific documents to train a topic-specific model. In our case the documents come from a first round of retrieval for the user’s current query, and the word embedding model is trained based on sentences from the topic-specific document set. Unlike the past work, we do not focus on interpolating the local and global models, although this is a promising area for future work. In the current study we focus on a direct comparison between the local-only and global-only approach, for improving retrieval performance.

Word embeddings for IR Information Retrieval has a long history of learning representations of words that are low-dimensional dense vectors. These approaches can be broadly classified into two families based on whether they are learnt based on a *term-document* matrix or *term co-occurrence* data. Using the *term-document* matrix for embedding leads to several well-studied approaches such as LSA (Deerwester et al., 1990), PLSA (Hofmann, 1999), and LDA (Blei et al., 2003; Wei and Croft, 2006). The performance of these models varies depending on the task, for example they are known to perform poorly for retrieval tasks unless combined with lexical features (Atreya and Elkan, 2011a). *Term-cooccurrence* based embeddings, such as word2vec (Mikolov et al., 2013b; Mikolov et al., 2013a) and (Pennington et al., 2014b), have recently been remarkably popular for many natural language processing and logical reasoning tasks. However, there are relatively less known successful applications of these models in IR. Ganguly *et. al.* (Ganguly et al., 2015) used the word similarity in the word2vec embedding space as a way to estimate term transformation probabilities in a language modelling setting for retrieval. More recently, Nalisnick *et. al.* (Nalisnick et al., 2016) proposed to model document about-ness by computing the similarity between all pairs of query and document terms using dual embedding spaces. Both these approaches estimate the semantic relatedness between two terms as the cosine distance between them in the embedding space(s). We adopt a similar notion of term relatedness but focus on demon-

strating improved retrieval performance using locally trained embeddings.

Local latent semantic analysis Despite the mathematical appeal of latent semantic analysis, several experiments suggest that its empirical performance may be no better than that of ranking using standard term vectors (Deerwester et al., 1990; Dumais, 1995; Atreya and Elkan, 2011b). In order to address the coarseness of corpus-level latent semantic analysis, Hull proposed restricting analysis to the documents relevant to a query (Hull, 1994). This approach significantly improved over corpus-level analysis for routing tasks, a result that has been reproduced in consequent research (Schütze et al., 1995; Singhal et al., 1997). Our work can be seen as an extension of these results to more recent techniques such as word2vec.

9 Conclusion

We have demonstrated a simple and effective method for performing query expansion with word embeddings. Importantly, our results highlight the value of locally-training word embeddings in a query-specific manner. The strength of these results suggests that other research adopting global embedding vectors should consider local embeddings as a potentially superior representation. Instead of using a “Sriracha sauce of deep learning,” as embedding techniques like word2vec have been called, we contend that the situation sometimes requires, say, that we make a béchamel or a mole verde or a sambal—or otherwise learn to cook.

References

Nasreen Abdul-Jaleel, James Allan, W. Bruce Croft, Fernando Diaz, Leah Larkey, Xiaoyan Li, Donald Metzler, Mark D. Smucker, Trevor Strohman, Howard Turtle, and Courtney Wade. 2004. Umass at trec 2004: Novelty and hard. In *Online Proceedings of 2004 Text REtrieval Conference*.

Mohannad al Masri, Catherine Berrut, and Jean-Pierre Chevallet. 2016. A comparison of deep learning based query expansion with pseudo-relevance feedback and mutual information. In Nicola Ferro, Fabio Crestani, Marie-Francine Moens, Josiane Mothe, Fabrizio Silvestri, Maria Giorgio Di Nunzio, Claudia Hauff,

and Gianmaria Silvello, editors, *Proceedings of the 38th European Conference on IR Research (ECIR 2016)*, pages 709–715, Cham. Springer International Publishing.

- Avinash Atreya and Charles Elkan. 2011a. Latent semantic indexing (lsi) fails for trec collections. *ACM SIGKDD Explorations Newsletter*, 12(2):5–10.
- Avinash Atreya and Charles Elkan. 2011b. Latent semantic indexing (lsi) fails for trec collections. *SIGKDD Explor. Newsl.*, 12(2):5–10, March.
- R. Attar and A. S. Fraenkel. 1977. Local feedback in full-text retrieval systems. *J. ACM*, 24(3):397–417, July.
- Jerome R Bellegarda. 2004. Statistical language model adaptation: review and perspectives. *Speech communication*, 42(1):93–108.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- W. Bruce Croft and John Lafferty. 2003. *Language Modeling for Information Retrieval*. Kluwer Academic Publishing.
- Steve Cronen-Townsend, Yun Zhou, and W. Bruce Croft. 2002. Predicting query performance. In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 299–306, New York, NY, USA. ACM Press.
- Scott C. Deerwester, Susan T. Dumais, Thomas K. Landauer, George W. Furnas, and Richard A. Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American Society of Information Science*, 41(6):391–407.
- Fernando Diaz and Donald Metzler. 2006. Improving the estimation of relevance models using large external corpora. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 154–161, New York, NY, USA. ACM Press.
- Fernando Diaz. 2015. Condensed list relevance models. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval, ICTIR '15*, pages 313–316, New York, NY, USA, May. ACM.
- Susan T. Dumais. 1995. Latent semantic indexing (LSI): TREC-3 report. In *Overview of the Third Text REtrieval Conference (TREC-3)*, pages 219–230.
- Marcello Federico. 1996. Bayesian estimation methods for n-gram language model adaptation. In *Spoken Language, 1996. ICSLP 96. Proceedings., Fourth International Conference on*, volume 1, pages 240–243. IEEE.

- William A. Gale, Kenneth W. Church, and David Yarowsky. 1992. One sense per discourse. In *Proceedings of the Workshop on Speech and Natural Language*, HLT '91, pages 233–237, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Debasis Ganguly, Dwaipayan Roy, Mandar Mitra, and Gareth J.F. Jones. 2015. Word embedding based generalized language model for information retrieval. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 795–798, New York, NY, USA. ACM.
- Zellig S. Harris. 1954. Distributional structure. *WORD*, 10(2-3):146–162.
- Thomas Hofmann. 1999. Probabilistic latent semantic indexing. In *SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57, New York, NY, USA. ACM Press.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 873–882. Association for Computational Linguistics.
- David Hull. 1994. Improving text retrieval for the routing problem using latent semantic indexing. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '94, pages 282–291, New York, NY, USA. Springer-Verlag New York, Inc.
- R.M. Iyer and M. Ostendorf. 1999. Modeling long distance dependence in language: topic mixtures versus dynamic cache models. *Speech and Audio Processing, IEEE Transactions on*, 7(1):30–39, Jan.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of ir techniques. *TOIS*, 20(4):422–446.
- Robert Krovetz. 1993. Viewing morphology as an inference process. In *SIGIR '93: Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 191–202, New York, NY, USA. ACM Press.
- Roland Kuhn and Renato De Mori. 1990. A cache-based natural language model for speech recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(6):570–583.
- Victor Lavrenko and W. Bruce Croft. 2001. Relevance based language models. In *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 120–127. ACM Press.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2177–2185. Curran Associates, Inc.
- Yang Liu, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2015. Topical word embeddings. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 2418–2424. AAAI Press.
- 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, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. 2016. Improving document ranking with dual word embeddings. In *Proc. WWW. International World Wide Web Conferences Steering Committee*.
- Hiroaki Nanjo and Tatsuya Kawahara. 2004. Language model and speaking rate adaptation for spontaneous presentation speech recognition. *Speech and Audio Processing, IEEE Transactions on*, 12(4):391–400.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2015. Efficient non-parametric estimation of multiple embeddings per word in vector space. *arXiv preprint arXiv:1504.06654*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014a. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014b. Glove: Global vectors for word representation. *Proc. EMNLP*, 12:1532–1543.
- Joseph Reisinger and Raymond Mooney. 2010a. A mixture model with sharing for lexical semantics. In *Proceedings of the 2010 Conference*

- on *Empirical Methods in Natural Language Processing*, pages 1173–1182. Association for Computational Linguistics.
- Joseph Reisinger and Raymond J Mooney. 2010b. Multi-prototype vector-space models of word meaning. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 109–117. Association for Computational Linguistics.
- Hinrich Schütze, David A. Hull, and Jan O. Pedersen. 1995. A comparison of classifiers and document representations for the routing problem. In *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '95, pages 229–237, New York, NY, USA. ACM.
- Hidetoshi Shimodaira. 2000. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227 – 244.
- Amit Singhal, Mandar Mitra, and Chris Buckley. 1997. Learning routing queries in a query zone. *SIGIR Forum*, 31(SI):25–32, July.
- Alessandro Sordoni, Yoshua Bengio, and Jian-Yun Nie. 2014. Learning concept embeddings for query expansion by quantum entropy minimization. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI'14, pages 1586–1592. AAAI Press.
- Anastasios Tombros and C. J. van Rijsbergen. 2001. Query-sensitive similarity measures for the calculation of interdocument relationships. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 17–24, New York, NY, USA. ACM Press.
- Anastasios Tombros, Robert Villa, and C. J. Van Rijsbergen. 2002. The effectiveness of query-specific hierarchic clustering in information retrieval. *Inf. Process. Manage.*, 38(4):559–582, July.
- Andrew Trask, Phil Michalak, and John Liu. 2015. sense2vec-a fast and accurate method for word sense disambiguation in neural word embeddings. *arXiv preprint arXiv:1511.06388*.
- Laurens van der Maaten and Geoffrey E. Hinton. 2008. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605.
- C. J. van Rijsbergen. 1979. *Information Retrieval*. Butterworths.
- Xing Wei and W. Bruce Croft. 2006. LDA-based document models for ad-hoc retrieval. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 178–185, New York, NY, USA. ACM Press.
- Peter Willett. 1985. Query-specific automatic document classification. In *International Forum on Information and Documentation*, volume 10, pages 28–32.
- Jinxi Xu and W. Bruce Croft. 1996. Query expansion using local and global document analysis. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '96, pages 4–11, New York, NY, USA. ACM.
- David Yarowsky. 1993. One sense per collocation. In *Proceedings of the Workshop on Human Language Technology*, HLT '93, pages 266–271, Stroudsburg, PA, USA. Association for Computational Linguistics.
- David Yarowsky. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics*, ACL '95, pages 189–196, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Bing Zhao, Matthias Eck, and Stephan Vogel. 2004. Language model adaptation for statistical machine translation with structured query models. In *Proceedings of the 20th International Conference on Computational Linguistics*, COLING '04, Stroudsburg, PA, USA. Association for Computational Linguistics.