

Neural Vector Spaces for Unsupervised Information Retrieval

CHRISTOPHE VAN GYSEL, MAARTEN DE RIJKE, and EVANGELOS KANOULAS,
University of Amsterdam, The Netherlands

We propose the Neural Vector Space Model (NVSM), a method that learns representations of documents in an unsupervised manner for news article retrieval. In the NVSM paradigm, we learn low-dimensional representations of words and documents from scratch using gradient descent and rank documents according to their similarity with query representations that are composed from word representations. We show that NVSM performs better at document ranking than existing latent semantic vector space methods. The addition of NVSM to a mixture of lexical language models and a state-of-the-art baseline vector space model yields a statistically significant increase in retrieval effectiveness. Consequently, NVSM adds a complementary relevance signal. Next to semantic matching, we find that NVSM performs well in cases where lexical matching is needed. NVSM learns a notion of term specificity directly from the document collection without feature engineering. We also show that NVSM learns regularities related to Luhn significance. Finally, we give advice on how to deploy NVSM in situations where model selection (e.g., cross-validation) is infeasible. We find that an unsupervised ensemble of multiple models trained with different hyperparameter values performs better than a single cross-validated model. Therefore, NVSM can safely be used for ranking documents without supervised relevance judgments.

CCS Concepts: • **Information systems** → **Content analysis and feature selection; Retrieval models and ranking**; *Document representation; Query representation*;

Additional Key Words and Phrases: Ad-hoc retrieval, latent vector spaces, semantic matching, document retrieval, representation learning

ACM Reference format:

Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2018. Neural Vector Spaces for Unsupervised Information Retrieval. *ACM Trans. Inf. Syst.* 36, 4, Article 38 (June 2018), 25 pages.

<https://doi.org/10.1145/3196826>

This research was supported by Ahold Delhaize, Amsterdam Data Science, the Bloomberg Research Grant program, the Criteo Faculty Research Award program, Elsevier, the European Community's Seventh Framework Programme (FP7/2007-2013) under Grant Agreement No. 312827 (VOX-Pol), the Google Faculty Research Award scheme, the Microsoft Research Ph.D. program, the Netherlands Institute for Sound and Vision, the Netherlands Organisation for Scientific Research (NWO) under Projects No. 612.001.116, No. CI-14-25, No. 652.002.001, No. 612.001.551, No. 652.001.003, and Yandex. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

Authors' addresses: C. Van Gysel, M. de Rijke, and E. Kanoulas, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands; emails: {cvangysel, derijke, e.kanoulas}@uva.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1046-8188/2018/06-ART38 \$15.00

<https://doi.org/10.1145/3196826>

1 INTRODUCTION

The vocabulary mismatch between query and document poses a critical challenge in search (Li and Xu 2014). The vocabulary gap occurs when documents and queries, represented as a bag-of-words, use different terms to describe the same concepts. While improved semantic matching methods are urgently needed, for these methods to be effective they need to be applicable at early stages of the retrieval pipeline. Otherwise, candidate documents most affected by the mismatch (i.e., relevant documents that do not contain any query terms) will simply remain undiscovered. Boytsov et al. (2016) show that (approximate) nearest neighbor algorithms (Garcia et al. 2008; Muja and Lowe 2014) can be more efficient than classical term-based retrieval. This strongly motivates the design of semantic matching methods that represent queries and documents in finite-dimensional vector spaces.

Latent semantic models, such as LSI (Deerwester et al. 1990), fit in the finite-dimensional vector space paradigm needed for nearest neighbor retrieval. However, LSI is known to retrieve non-relevant documents due to a lack of specificity (Dumais 1995). The recent move toward learning word representations as part of neural language models (Bengio et al. 2003) has shown impressive improvements in natural language processing (NLP) (Collobert et al. 2011; Graves and Jaitly 2014; Mikolov et al. 2013b). Therefore, it is reasonable to explore these representation learning methods for information retrieval (IR) as well. Unfortunately, in the case of full text document retrieval, only few positive results have been obtained so far (Craswell et al. 2016). We identify two causes for this shortfall.

First of all, IR tasks (e.g., document ranking) are fundamentally different from NLP tasks (Craswell et al. 2016). NLP deals with natural language regularities (e.g., discovering long range dependencies), whereas IR involves satisfying a user's information need (e.g., matching a query to a document). Therefore, specialized solutions and architectures for IR are needed.

Second, in the bag-of-words paradigm, query/document matching is performed by counting term occurrences within queries and documents. Afterward, the frequencies are adjusted using weighting schemes that favor term specificity (Robertson 2004; Robertson and Walker 1994), used as parameter values in probabilistic frameworks (Zhai and Lafferty 2004) and/or reduced in dimensionality (Blei et al. 2003; Deerwester et al. 1990). However, Baroni et al. (2014) noted that for NLP tasks, prediction-based models, learned from scratch using gradient descent, outperform count-based models. Similarly, the advances made by deep learning in computer vision (Krizhevsky et al. 2012) were not due to counting visual words (Vidal-Naquet and Ullman 2003) or by constructing complicated pipelines, but instead by optimizing a cost function using gradient descent and learning the model from scratch. However, for IR settings such as unsupervised news article retrieval, where one ranks documents in the absence of explicit or implicit relevance labels, prediction-based models have not received much attention. For deep learning to become feasible for unsupervised retrieval, we first need to construct an appropriate optimization objective.

In this article, we introduce an optimization objective for learning latent representations of words and documents from scratch, in an unsupervised manner without relevance signals. Specifically, we introduce the Neural Vector Space Model (NVSM) for document retrieval. The optimization objective of NVSM mandates that word sequences extracted from a document should be predictive of that document. Learning a model of content in this way incorporates the following IR regularities:

- *Semantic Matching*. Words occurring in each other's vicinity are learned to have similar representations,
- *The Clustering Hypothesis*. Documents that contain similar language will have nearby representations in latent space.

- *Term Specificity*. Words associated with many documents will be neglected, due to low predictive power.

One limitation of latent document vector spaces, including the NVSM we introduce here, is that their asymptotic complexity is bounded by the number of documents (i.e., one vector for every document) (Ai et al. 2016b; Chen 2017; Van Gysel et al. 2016a). Consequently, the latent methods we consider in this article are only feasible to be constructed on document collections of medium scale. Therefore, we choose to evaluate our methods on article retrieval benchmarks (~200k to 2m documents each) from TREC (Harman 1992; Zhai and Lafferty 2004).

We seek to answer the following research questions:

- (1) How does NVSM compare to other latent vector space models, such as doc2vec (Le and Mikolov 2014), word2vec (Mikolov et al. 2013b; Vulić and Moens 2015), LSI (Deerwester et al. 1990), LDA (Blei et al. 2003), and LSE (Van Gysel et al. 2016a), on the document retrieval task?
- (2) For what proportion of queries does NVSM perform better than the other rankers?
- (3) What gains does NVSM bring when combined with a lexical Query Likelihood Model (QLM) and a competing state-of-the-art vector space model?
- (4) Do NVSMs exhibit regularities that we can link back to well-understood document collection statistics used in traditional retrieval models and how do the regularities in NVSM differ from those in LSE?

We contribute:

- (1) A novel neural retrieval model, NVSM, that is learned using gradient descent on a document collection.
- (2) Comparisons of lexical (QLM) and semantic (doc2vec, word2vec, LSI, LDA, LSE, and NVSM) models on document retrieval test collections.
- (3) An analysis of the internals of NVSM to give insights in the workings of the model and the retrieval task.
- (4) A highly-optimized open source C++/CUDA implementation of NVSM that results in fast training and efficient memory usage.¹
- (5) Advice on how to configure the hyperparameters of NVSM.

2 RELATED WORK

The majority of retrieval models represent documents and queries as elements of a finite-dimensional vector space and rank documents according to some similarity measure defined in this space. A distinction is made between bag-of-words models (Robertson and Walker 1994; Sparck Jones 1972; Zhai and Lafferty 2004), for which the components of the vector space correspond to terms, and latent models (Ai et al. 2016b; Blei et al. 2003; Deerwester et al. 1990; Hofmann 1999; Van Gysel et al. 2016a), where the components are unobserved and the space is of low dimensionality compared to the number of unique terms. We first discuss prior work on representation learning for document retrieval, with an emphasis on neural methods. Next, we give an overview of neural language modeling for NLP and automatic speech recognition.

2.1 Unsupervised Representations for IR

The recent revival of neural networks due to advances in computer vision (Krizhevsky et al. 2012), NLP (Collobert et al. 2011; Mikolov et al. 2013b), and automatic speech processing (ASR) (Graves

¹<https://github.com/cvangysel/cuNVSM>.

and Jaitly 2014) have led to an increasing interest in these technologies from the IR community. In particular, word embeddings (Bengio et al. 2003), low-dimensional representations of words learned as part of a neural language model, have attracted much attention due to their ability to encode semantic relations between words. First we discuss non-neural latent semantic models, followed by work that integrates pre-trained neural embeddings into existing retrieval models, and work on neural retrieval models that are learned from scratch.

2.1.1 Latent Semantic Models. Latent Semantic Indexing (LSI) (Deerwester et al. 1990) and probabilistic LSI (pLSI) (Hofmann 1999) were introduced to mitigate the mismatch between documents and queries (Li and Xu 2014). Blei et al. (2003) proposed Latent Dirichlet Allocation (LDA), a topic model that generalizes to unseen documents.

2.1.2 Combining Pre-Trained Word or Document Representations. The following methods take pre-existing representations and use them to either:

- (1) Obtain a document representation from individual word representations that are subsequently used for (re-)ranking, or
- (2) Combine representation similarities in some way to (re-)rank documents.

While many of these methods were applied to word2vec (Mikolov et al. 2013b) representations at the time of publishing, these methods are in fact generic in the sense that they can be applied to any model that provides representations of the required unit (i.e., words or documents). Consequently, most of the methods below can be applied to NVSM and the baseline vector space models we consider in this article. However, the key focus of this article is the alleviation of the vocabulary gap in IR, and, consequently, for practical reasons, nearest neighbor search is the most viable way to compare latent vector space models (Boytsov et al. 2016).

Vulić and Moens (2015) are the first to aggregate word embeddings learned with a context-predicting distributional semantic model (DSM); query and document are represented as a sum of word embeddings learned from a pseudo-bilingual document collection with a Skip-Gram model. Kenter and de Rijke (2015) extract features from embeddings for the task of determining short text similarity. Zuccon et al. (2015) use embeddings to estimate probabilities in a translation model that is combined with traditional retrieval models (similar to Ganguly et al. (2015) and Tu et al. (2016)). Zamani and Croft (2016a, 2016b) investigate the use of pre-trained word embeddings for query expansion and as a relevance model to improve retrieval. Guo et al. (2016b) introduce the Bag-of-Word-Embeddings (BoWE) representation where every document is represented as a matrix of the embeddings occurring in the document; their non-linear word transportation model compares all combinations of query/document term representations at retrieval time. They incorporate lexical matching into their model by exactly comparing embedding vector components for specific terms (i.e., specific terms occurring in both document and query are matched based on the equality of their vector components, contrary to their lexical identity).

2.1.3 Learning from Scratch. The methods discussed above incorporate features from neural language models. The recent deep learning revival, however, was due to the end-to-end optimization of objectives and representation learning (Krizhevsky et al. 2012; LeCun et al. 1998; Sutskever et al. 2014) in contrast to feature engineering or the stacking of independently-estimated models. The following neural methods learn representations of words and documents from scratch. Salakhutdinov and Hinton (2009) introduce semantic hashing for the document similarity task. Le and Mikolov (2014) propose doc2vec, a method that learns representations of words and documents. Ai et al. (2016b) evaluate the effectiveness of doc2vec representations for ad hoc retrieval, but obtain disappointing results that are further analyzed in (Ai et al. 2016a). An important

difference between the work of Ai et al. and this article is that we operate within the vector space framework. This allows us to retrieve documents using an (approximate) nearest neighbor search and thus mitigate the vocabulary gap at early stages of the retrieval pipeline. In contrast, the approach by Ai et al. (2016b) can only be feasibly applied in a re-ranking setting. Van Gysel et al. develop end-to-end neural language models for entity ranking tasks (Van Gysel et al. 2017b); they show how their neural models outperform traditional language models on the expert finding task (Van Gysel et al. 2016b) and improve the scalability of their approach and extend it to product search (Van Gysel et al. 2016a).

2.2 Neural IR and Language Modeling

Beyond representation learning, there are more applications of neural models in IR (Craswell et al. 2016; Kenter et al. 2017; Onal et al. 2018; Van Gysel et al. 2017c). In machine-learned ranking (Liu 2011), we have RankNet (Burgess et al. 2005). In the class of supervised learning-to-match approaches, where clicks are available, there are DSSM (Huang et al. 2013; Shen et al. 2014) and DSN (Deng et al. 2013). Guo et al. (2016a) learn a relevance model by extracting features from BoWE representations in addition to corpus statistics such as inverse document frequency (IDF). Recently, Mitra et al. (2017) have introduced a supervised document ranking model that matches using both local and distributed representations. Next to retrieval models, there has been work on modeling user interactions with neural methods. Borisov et al. (2016b) introduce a neural click model that represents user interactions as a vector representation; in Borisov et al. (2016a), they extend their work by taking into account click dwell time.

Neural Network Language Models (NNLM) (Bengio et al. 2003) have shown promising results in NLP (Jozefowicz et al. 2015; Sutskever et al. 2014; Tran et al. 2016) and ASR (Graves and Jaitly 2014; Sak et al. 2014) compared to Markovian models. Collobert et al. (2011) apply NNLMs to arbitrary NLP tasks by learning one set of word representations in a multi-task setting. Turian et al. (2010) compare word representations learned by neural networks, distributional semantics, and cluster-based methods as features for NLP tasks (Baroni et al. 2014). Following this observation, models have been designed with the sole purpose of learning embeddings (Kenter et al. 2016; Mikolov et al. 2013b; Pennington et al. 2014). Levy and Goldberg (2014) show the relation between neural word embeddings and distributional word representations and how to improve the latter using lessons learned from the former (Levy et al. 2015).

The contribution of this article over and above the related work discussed above is the following. First, we substantially extend the LSE model of Van Gysel et al. (2016a) and evaluate it on document search collections from TREC (TREC 1999). Our improvements over the LSE model are due to:

- (1) Increased regularization
- (2) Accelerated training by reducing the internal covariate shift
- (3) The incorporation of term specificity within the learned word representations

Second, we avoid *irrational exuberance* known to plague the deep learning field (Craswell et al. 2016) by steering away from non-essential depth. That is, while we extend algorithms and techniques from deep learning, the model we present in this article is shallow.

Third, and contrary to previous work where information from pre-trained word embeddings is used to enrich the query/document representation of count-based models, our model, NVSM, is learned directly from the document collection without explicit feature engineering. We show that NVSM learns regularities known to be important for document retrieval from scratch. Our semantic vector space outperforms lexical retrieval models on some benchmarks. However, given that lexical and semantic models perform different types of matching, our approach is most useful as a supplementary signal to these lexical models.

3 LEARNING SEMANTIC SPACES

In this section, we provide the details of NVSM. First, we give a birds-eye overview of the model and its parameters and explain how to rank documents for a given query. Secondly, we outline our training procedure and optimization objective. We explain the aspects of the objective that make it work in a retrieval setting. Finally, we go into the technical challenges that arise when implementing the model and how we solved them in our open-source release.

3.1 The Neural Vector Space Model

Our work revolves around unsupervised ad hoc document retrieval where a user wishes to retrieve documents (e.g., articles) as to satisfy an information need encoded in query q .

Below, a query q consists of terms (i.e., words) $t_1, \dots, t_{|q|}$ originating from a vocabulary V , where $|\cdot|$ denotes the length operator; D denotes the set of documents $\{d_1, \dots, d_{|D|}\}$. Every document $d_i \in D$ consists of a sequence of words $w_{i,1}, \dots, w_{i,|d_i|}$ with $w_{i,j} \in V$, for all $1 \leq j \leq |d_i|$.

We extend the work of Van Gysel et al. (2016a) who learn low-dimensional representations of words, documents, and the transformation between them from scratch. That is, instead of counting term frequencies for use in probabilistic frameworks (Zhai and Lafferty 2004) or applying dimensionality reduction to term frequency vectors (Deerwester et al. 1990; Wei and Croft 2006), we learn representations directly by gradient descent from sampled n-gram/document pairs extracted from the corpus.² These representations are embedded parameters in matrices $R_D \in \mathbb{R}^{|D| \times k_d}$ and $R_V \in \mathbb{R}^{|V| \times k_w}$ for documents D and vocabulary words V , respectively, such that $\vec{R}_V^{(i)}$ ($\vec{R}_D^{(j)}$, respectively) denotes the k_w -dimensional (k_d -dimensional, respectively) vector representation of word V_i (document d_i , respectively).

As the word representations $\vec{R}_V^{(i)}$ and document representations $\vec{R}_D^{(j)}$ are of different dimensionality, we require a transformation $f: \mathbb{R}^{k_w} \rightarrow \mathbb{R}^{k_d}$ from the word feature space to the document feature space. In this article, we take the transformation to be linear:

$$f(\vec{x}) = W\vec{x}, \quad (1)$$

where \vec{x} is a k_w -dimensional vector and W is a $k_d \times k_w$ parameter matrix that is learned using gradient descent (in addition to representation matrices R_V and R_D).

We compose a representation of a sequence of n words (i.e., an n-gram) w_1, \dots, w_n by averaging its constituent word representations:

$$g(w_1, \dots, w_n) = \frac{1}{n} \sum_{i=1}^n \vec{R}_V^{(w_i)}. \quad (2)$$

A query q is projected to the document feature space by the composition of f and g : $(f \circ g)(q) = h(q)$.

The matching score between a document d and query q is then given by the cosine similarity between their representations in document feature space:

$$\text{score}(q, d) = \frac{h(q)^T \cdot \vec{R}_D^{(d)}}{\|h(q)\| \|\vec{R}_D^{(d)}\|}. \quad (3)$$

We then proceed by ranking documents $d \in D$ in decreasing order of score (q, d) (see Equation (3)) for a given query q . Note that ranking according to cosine similarity is equivalent to ranking

²Note here that the intention is not to use n-grams as a substitute for queries, but rather that n-grams (i.e., multiple words) contain more meaning (i.e., semantics) than a single word. The training procedure can thus be characterized as projecting the semantic representation of an n-gram close to the document from which the n-gram was extracted.

according to inverse Euclidean distance if the vectors are normalized. Therefore, ranking according to Equation (3) can be formulated as an (approximate) nearest neighbor search problem in a metric space.

The model proposed here, NVSM, is an extension of the LSE model (Van Gysel et al. 2016a). NVSM/LSE are different from existing unsupervised neural retrieval models learned from scratch due to their ability to scale to collections larger than expert finding collections (Van Gysel et al. 2016b) (i.e., $\sim 10k$ entities/documents) and the assumption that words and documents are embedded in different spaces (Le and Mikolov 2014). Compared to word embeddings (Mikolov et al. 2013b; Pennington et al. 2014), NVSM/LSE learn document-specific representations instead of collection-wide representations of language; see (Van Gysel et al. 2016a, p. 4) for a more in-depth discussion. The algorithmic contribution of NVSM over LSE comes from improvements in the objective that we introduce in the next section.

3.2 The Objective and Its Optimization

We learn representations of words and documents using mini-batches of m n-gram/document pairs such that an n-gram representation—composed out of word representations—is projected nearby the document that contains the n-gram, similar to LSE (Van Gysel et al. 2016a). The word and n-gram representations, and the transformation between them are all learned simultaneously. This is in contrast to representing documents by a weighted sum of pre-trained representations of the words that it contains (Vulić and Moens 2015). A mini-batch B is constructed as follows:

- (1) Stochastically sample document $d \in D$ according to $P(D)$. In this article, we assume $P(D)$ to be uniform, similar to (Zhai and Lafferty 2004). Note that $P(D)$ can be used to incorporate importance-based information (e.g., document length).
- (2) Sample a phrase of n contiguous words w_1, \dots, w_n from document d .
- (3) Add the phrase-document pair $(w_1, \dots, w_n; d)$ to the batch.
- (4) Repeat until the batch is full.

Given a batch B , we proceed by constructing a differentiable, non-convex loss function L of the parameters Θ (e.g., R_V, R_D) and the parameter estimation problem is cast as an optimization problem that we approximate using stochastic gradient descent.

Denote B_i as the i -th n-gram/document pair of batch B and $B_i^{(p)}$ ($B_i^{(d)}$, respectively) as the n-gram (document, respectively) of pair B_i . Further, we introduce an auxiliary function that L2-normalizes a vector of arbitrary dimensionality:

$$\text{norm}(\vec{x}) = \frac{\vec{x}}{\|\vec{x}\|}.$$

For an n-gram/document pair B_i , the non-standardized projection of the n-gram into the k_d -dimensional document feature space is as follows:

$$\tilde{T}(B_i^{(p)}) = (f \circ \text{norm} \circ g)(B_i^{(p)}). \quad (4)$$

A few quick comments are in order. The function g (see Equation (2)) constructs an n-gram representation by averaging the word representations (embedded in the R_V parameter matrix). This allows the model to learn semantic relations between words for the purpose of *semantic matching*. That is, the model does not learn from individual words, but instead it learns from an unordered sequence (order is not preserved as we sum in Equation (2)) of words that constitute meaning. As documents contain multiple n-grams and n-grams are made up of multiple words, semantic similarity between words and documents is learned. In addition, the composition function g in

combination with L2-normalization norm (\cdot) causes words to compete to contribute to the resulting n-gram representation. Given that we will optimize the n-gram representation to be close to the corresponding document (as we will explain below), words that are discriminative for the document in question will learn to contribute more to the n-gram representation (due to their discriminative power), and consequently, the L2-norm of the representations of discriminative words will be larger than the L2-norm of non-discriminative words. This incorporates a notion of *term specificity* into our model.

We then estimate the per-feature sample mean and variance

$$\hat{\mathbb{E}} [\tilde{T} (B_i^{(p)})] \text{ and } \hat{\mathbb{V}} [\tilde{T} (B_i^{(p)})]$$

over batch B . The standardized projection of n-gram/document pair B_i can then be obtained as follows:

$$T (B_i^{(p)}) = \text{hard-tanh} \left(\frac{\tilde{T} (B_i^{(p)}) - \hat{\mathbb{E}} [\tilde{T} (B_i^{(p)})]}{\sqrt{\hat{\mathbb{V}} [\tilde{T} (B_i^{(p)})]} + \beta} \right), \quad (5)$$

where β is a k_d -dimensional bias vector parameter that captures document-independent regularities corresponding to word frequency. While vector β is learned during training, it is ignored during prediction (i.e., a nuisance parameter) similar to the position bias in click models (Chapelle and Zhang 2009) and score bias in learning-to-rank (Joachims 2002). The standardization operation reduces the internal covariate shift (Ioffe and Szegedy 2015). That is, it avoids the complications introduced by changes in the distribution of document feature vectors during learning. In addition, as the document feature vectors are learned from scratch as well, the standardization forces the learned feature vectors to be centered around the null vector. Afterward, we apply the hard-saturating nonlinearity hard-tanh (Gulcehre et al. 2016) such that the feature activations are between -1 and 1 .

The objective is to maximize the similarity between

$$T (B_i^{(p)}) \text{ and } \vec{R}_D^{(B_i^{(d)})},$$

while minimizing the similarity between

$$T (B_i^{(p)})$$

and the representations of other documents. Therefore, a document is characterized by the concepts it contains, and, consequently, documents describing similar concepts will cluster together, as postulated by the *clustering hypothesis* (van Rijsbergen 1979). NVSM strongly relies on the clustering hypothesis as ranking is performed according to nearest neighbor retrieval (Equation (3)).

Considering the full set of documents D is often costly as $|D|$ can be large. Therefore, we apply an adjusted-for-bias variant of negative sampling (Gutmann and Hyvärinen 2010; Mikolov et al. 2013a; Van Gysel et al. 2016a), where we uniformly sample negative examples from D . Adopted from (Van Gysel et al. 2016a), we define

$$P (\mathcal{S} \mid d, B_i^{(p)}) = \sigma (\vec{R}_D^{(d)} \cdot T (B_i^{(p)})) \quad (6)$$

as the similarity of two representations in latent vector space, where

$$\sigma (t) = \frac{1}{1 + \exp (-t)}$$

denotes the sigmoid function and \mathcal{S} is an indicator binary random variable that says whether the representation of document d is similar to the projection of n-gram $B_i^{(p)}$.

The probability of document $B_i^{(d)}$ given phrase $B_i^{(p)}$ is then approximated by uniformly sampling z contrastive examples:

$$\log \tilde{P} \left(B_i^{(d)} \mid B_i^{(p)} \right) = \frac{z+1}{2z} \left(z \log P \left(\mathcal{S} \mid B_i^{(d)}, B_i^{(p)} \right) + \sum_{\substack{k=1, \\ d_k \sim U(D)}}^z \log \left(1.0 - P \left(\mathcal{S} \mid d_k, B_i^{(p)} \right) \right) \right), \quad (7)$$

where $U(D)$ denotes the uniform distribution over documents D , the distribution used for obtaining negative examples (Gutmann and Hyvärinen 2010; Van Gysel et al. 2016a). Then, the *loss function* we use to optimize our model is Equation (7) averaged over the instances in batch B :

$$L(R_V, R_D, W, \beta \mid B) = -\frac{1}{m} \sum_{i=1}^m \log \tilde{P} \left(B_i^{(d)} \mid B_i^{(p)} \right) + \frac{\lambda}{2m} \left(\sum_{i,j} R_V^2_{i,j} + \sum_{i,j} R_D^2_{i,j} + \sum_{i,j} W^2_{i,j} \right), \quad (8)$$

where λ is a weight regularization hyper-parameter. We optimize our parameters θ (R_V , R_D , W and β) using Adam (Kingma and Ba 2014), a first-order gradient-based optimization function for stochastic objective functions that is very similar to momentum. The update rule for parameter θ given a batch B at batch update time t equals

$$\theta^{(t+1)} \leftarrow \theta^{(t)} - \alpha \frac{\hat{m}_\theta^{(t)}}{\sqrt{\hat{v}_\theta^{(t)} + \epsilon}}, \quad (9)$$

where $\hat{m}_\theta^{(t)}$ and $\hat{v}_\theta^{(t)}$, respectively, are the first- and second-moment estimates (over batch update times) (Kingma and Ba 2014) of the gradient of the loss $\frac{\partial L}{\partial \theta} (R_V, R_D, W, \beta \mid B)$ w.r.t. parameter θ at batch update time t and $\epsilon = 10^{-8}$ is a constant to ensure numerical stability. The use of this optimization method causes every parameter to be updated with every batch, unlike regular stochastic gradient descent, where the only parameters that are updated are those for which there is a non-zero gradient of the loss. This is important in NVSM due to the large number of word and document vectors.

The algorithmic contributions of NVSM over LSE are the components of the objective mentioned next. Equation (4) forces individual words to compete to contribute to the resulting n-gram representation. Consequently, non-discriminative words will have a small L2-norm. In Equation (5), we perform standardization to reduce the internal covariate shift (Ioffe and Szegedy 2015). In addition, the standardization forces n-gram representations to distinguish themselves only in the dimensions that matter for matching. Frequent words are naturally prevalent in n-grams, however, they have low discriminative power as they are non-specific. The bias β captures word frequency regularities that are non-discriminative for the semantic concepts within the respective n-gram/document pairs and allows the transformation in Equation (1) to focus on concept matching. The re-weighting of the positive instance in Equation (7) removes a dependence on the number of negative examples z , where a large z presented the model with a bias towards negative examples.

3.3 Implementation

The major cause of technical challenges of the NVSM training procedure is not due to time complexity, but rather space restrictions. This is because we mitigate expensive computation by estimating vector space models using graphics processing units (GPUs). The main limitation of these massively parallel computation devices is that they rely on their own memory units. Consequently, parameters and intermediate results of the training procedure need to persist in the GPU memory space. The asymptotic space complexity of the parameters equals

$$O \left(\underbrace{|V| \times k_w}_{\text{word representations } R_V} + \underbrace{k_d \times k_w}_{\text{transform } W} + \underbrace{|D| \times k_d}_{\text{document representations } R_D} \right).$$

In addition, Equation (9) requires us to keep the first and second moment of the gradient over time for every parameter in memory. Therefore, for every parameter, we retain three floating point values in memory at all times. For example, if we have a collection of 1M documents (256-dim.) with a vocabulary of 64K terms (300-dim.), then the model has ~275M parameters. Consequently, under the assumption of 32-bit floating point, the resident memory set has a best-case least upper bound of 3.30GB memory. The scalability of our method—similar to all latent vector space models—is determined by the number of documents within the collection. However, the current generation of GPUs—that typically boast around 12GB of memory—can comfortably be used to train models of collections consisting of up to 2 million documents. In fact, next-generation GPUs have double the amount of memory—24GB—and this amount will likely increase with the introduction of future processing units (Wikipedia 2017). This, and the development of distributed GPU technology (Abadi et al. 2015), leads us to believe that the applicability of our approach to larger retrieval domains is simply a matter of time (Moore 1998).

In addition to the scarcity of memory on current generation GPUs, operations such as the averaging of word to n-gram representations (Equation (2)) can be performed in-place. However, these critical optimizations are not available in general-purpose machine learning toolkits. Therefore, we have implemented the NVSM training procedure directly in C++/CUDA, such that we can make efficient use of sparseness and avoid unnecessary memory usage. In addition, models are stored in the open HDF5 format (Folk et al. 2011) and the toolkit provides a Python module that can be used to query trained NVSM models on the CPU. This way, a trained NVSM can easily be integrated in existing applications. The toolkit is licensed under the permissive MIT open-source license; see footnote 1.

4 EXPERIMENTAL SETUP

4.1 Research Questions

In this article, we investigate the viability of neural representation learning methods for semantic matching in document search. We seek to answer the following research questions:

RQ1 How does NVSM compare to other latent vector space models, such as doc2vec (Le and Mikolov 2014), word2vec (Mikolov et al. 2013b; Vulić and Moens 2015), LSI (Deerwester et al. 1990), LDA (Blei et al. 2003) and LSE (Van Gysel et al. 2016a), on the document retrieval task?

In particular, how does it compare to other methods that represent queries/documents as low-dimensional vectors? What is the difference in performance with purely lexical models that perform exact term matching and represent queries/documents as bag-of-words vectors?

RQ2 For what proportion of queries does NVSM perform better than the other rankers?

Does NVSM improve over other retrieval models only on a handful of queries? Instead of computing averages over queries, what if we look at the pairwise differences between rankers?

RQ3 What gains does NVSM bring when combined with a lexical QLM and a competing state-of-the-art vector space model?

Can we use the differences that we observe in **RQ2** between NVSM, QLM and other latent vector space models to our advantage to improve retrieval performance? Can we pinpoint where the improvements of NVSM come from?

RQ4 Do NVSMs exhibit regularities that we can link back to well-understood document collection statistics used in traditional retrieval models and how do the regularities in NVSM differ from those in LSE?

Traditional retrieval models such as generative language models are known to incorporate corpus statistics regarding term specificity (Robertson 2004; Sparck Jones 1972) and document length (Zhai and Lafferty 2004). Are similar corpus statistics present in NVSM and what does this tell us about the ranking task?

4.2 Benchmark Datasets & Experiments

In this article, we are interested in query/document matching. Therefore, we evaluate NVSM on newswire article collections from TREC. Other retrieval domains, such as web search or social media, deal with aspects such as document freshness/importance and social/hyperlink/click graphs that may obfuscate the impact of matching queries to documents. Therefore, we follow the experimental setup of Zhai and Lafferty (2004) and use four article retrieval sub-collections from the TIPSTER corpus (Harman 1992): Associated Press 88-89 (AP88-89), *Financial Times* (FT), *LA Times* (LA) and *Wall Street Journal* (WSJ) (Harman 1993). In addition, we consider the Robust04 collection that constitutes Disk 4/5 of the TIPSTER corpus without the Congressional Record and the *New York Times* (NY) collection that consists of articles written and published by the *New York Times* between 1987 and 2007. For evaluation, we take topics 50–200 from TREC 1–3³ (AP88-89, WSJ), topics 301–450 from TREC 6–8³ (FT, LA) (TREC 1999), topics 301–450, 601–700 from Robust04 (Voorhees 2005) and the 50 topics assessed by NIST (a subset of the Robust04 topics judged for the NY collection) during the TREC 2017 Common Core Track (Allan et al. 2017). From each topic, we take its title as the corresponding query. Topics without relevant documents are filtered out. We randomly create a split⁴ of validation (20%) and test (80%) queries (with the exception of the NY collection). For the NY collection, we select the hyperparameter configuration that optimizes the Robust04 validation set on the Robust04 collection and we take the 50 queries assessed by NIST and their judgments—specifically created for the NY collection—as our test set. This way, method hyperparameters are optimized on the validation set (as described in Section 4.3) and the retrieval performance is reported on the test set; see Table 1.

The inclusion of early TREC collections (AP88-89, FT, LA, and WSJ) is motivated by the fact that during the first few years of TREC, there was a big emphasis on submissions where the query was constructed manually from each topic and interactive feedback was used (Harman 1993). That is, domain experts repeatedly formulated manual queries using the full topic (title, description, narrative), observed the obtained rankings, and then reformulated their query to obtain better rankings. Consequently, these test collections are very useful when evaluating semantic matches as relevant documents do not necessarily contain topic title terms. From TREC-5 and onward, less emphasis

³We only consider judgments corresponding to each of the sub-collections.

⁴The validation/test splits can be found at <https://github.com/cvangysel/cuNVSM>.

Table 1. Overview of the Retrieval Benchmarks

	AP88-89	FT	LT
Collection (training)			
Documents	164,597	210,158	131,896
Document length	461.63 ± 243.61	399.68 ± 366.41	502.18 ± 519.58
Unique terms	2.67 × 10 ⁵	3.05 × 10 ⁵	2.67 × 10 ⁵
Queries (testing)			
	(T) 119	(T) 116	(T) 113
	(V) 30	(V) 28	(V) 30
Query terms	5.06 ± 3.14	2.50 ± 0.69	2.48 ± 0.69
Relevant documents	(T) 111.45 ± 136.01 (V) 86.43 ± 72.63	(T) 34.91 ± 42.94 (V) 30.46 ± 26.97	(T) 24.83 ± 34.31 (V) 24.30 ± 21.19
	NY	Robust04	WSJ
Collection (training)			
Documents	1,855,658	528,155	173,252
Document length	572.18 ± 605.82	479.72 ± 869.27	447.51 ± 454.75
Unique terms	1.35 × 10 ⁶	7.83 × 10 ⁵	2.50 × 10 ⁵
Queries (testing)			
	(T) 50	(T) 200	(T) 120
		(V) 49	(V) 30
Query terms	6.58 ± 0.70	5.28 ± 0.74	5.05 ± 3.14
Relevant documents	(T) 180.04 ± 132.74	(T) 70.33 ± 73.68 (V) 68.27 ± 77.41	(T) 96.99 ± 93.30 (V) 101.93 ± 117.65

T and V denote the test and validation sets, respectively. Arithmetic mean and standard deviation are reported wherever applicable.

was put on rankings generated using interactive feedback and shifted toward automated systems only (Harman and Voorhees 1996, footnote 1). In fact, within the Robust04 track, only automated systems were submitted (Voorhees 2005, §2) due to the large number of (a) documents (~500K) and (b) topics (250). In the 2017 Common Core Track (Allan et al. 2017), interactive rankings were once again submitted by participants, in addition to rankings obtained by the latent vector space model presented in this article.

We address **RQ1** by comparing NVSM to latent retrieval models (detailed in Section 4.3). In addition, we perform a per-query pairwise comparison of methods where we look at what method performs best for each query in terms of MAP@1000 (**RQ2**). A method performs better or worse than another if the absolute difference in MAP@1000 exceeds $\delta = 0.01$; otherwise, the two methods perform similar. To address **RQ3**, we consider the combinations (Section 4.3.2) of QLM with NVSM and the strongest latent vector space baseline of **RQ1**. That is, word2vec where the summands are weighted using self-information. In addition, we look at the correlation between per-query `titlestat_rel` (see Section 4.4) and the pairwise differences in MAP@1000 between NVSM and all the other retrieval models. A positive correlation indicates that NVSM is better at lexical matching than the other method, and vice versa for a negative correlation. For **RQ4**, we examine the relation between the collection frequency CF_w and the L2-norm of their word embeddings $g(w)$ for all terms $w \in V$.

4.3 Retrieval Models Considered for Comparison

The document collection is first indexed by Indri⁵ (Strohman et al. 2005). Retrieval models not implemented by Indri access the underlying tokenized document collection using `pyndri` (Van

⁵Stopwords are removed using the standard stopword list of Indri.

Gysel et al. 2017a). This way, all methods compared in this article parse the text collection consistently.

4.3.1 Models Compared. The key focus of this article is the alleviation of the vocabulary gap in IR and, consequently, in theory, we score all documents in each collection for every query. In practice, however, we rely on nearest neighbor search algorithms to retrieve the top-k documents (Boysov et al. 2016). Note that this is in contrast to many other semantic matching methods (Ai et al. 2016b; Nalisnick et al. 2016; Zuccon et al. 2015) that have only been shown to perform well in document re-ranking scenarios where an initial pool of candidate documents is retrieved using a lexical matching method. However, candidate documents most affected by the vocabulary gap (i.e., relevant documents that do not contain any query terms) will simply remain undiscovered in a re-ranking scenario and, consequently, we compare NVSM only to latent vector space models that can be queried using a nearest neighbor search.

The following latent vector space models are compared:

- (1) doc2vec (d2v) (Le and Mikolov 2014) with the distributed memory architecture. The pre-processing of document texts to learn latent document representations is a topic of study by itself and its effects are outside the scope of this article. Consequently, we disable vocabulary filtering and frequent word subsampling to keep the input to all representation learning algorithms consistent. We sweep the one-sided window size and the embedding size respectively in partitions $\{x/2 \mid x = 4, 6, 8, 10, 12, 16, 24, 32\}$ and $\{64, 128, 256\}$ on the validation set. Models are trained for 15 iterations and we select the model iteration that performs best on the validation set. Documents are ranked in decreasing order of the cosine similarity between the document representation and the average of the word embeddings in the query.
- (2) word2vec (w2v) (Mikolov et al. 2013b; Vulić and Moens 2015) with the Skip-Gram architecture. We follow the method introduced by Vulić and Moens (2015) where query/document representations are constructed by composing the representations of the words contained within them. We consider both the unweighted sum (add) and the sum of vectors weighted by the term's self-information (si). Self-information is a term specificity measure similar to IDF (Cover and Thomas 2012). The hyperparameters of word2vec are swept in the same manner as doc2vec.
- (3) LSI (Deerwester et al. 1990) with TF-IDF weighting and the number of topics $K \in \{64, 128, 256\}$ optimized on the validation set.
- (4) LDA (Blei et al. 2003) with $\alpha = \beta = 0.1$ and the number of topics $K \in \{64, 128, 256\}$ optimized on the validation set. We train the model for 100 iterations or until topic convergence is achieved. Documents are ranked in decreasing order of the cosine similarity between the query topic distribution and the document topic distribution.
- (5) Representation learning methods LSE (Van Gysel et al. 2016a) and NVSM (this article). For hyperparameters, we largely follow the findings of (Van Gysel et al. 2016a): word representation dim. $k_w = 300$, number of negative examples $z = 10$, learning rate $\alpha = 0.001$, regularization lambda $\lambda = 0.01$. For LSE, batch size $m = 4096$ (as in Van Gysel et al. (2016a)), while for NVSM the batch size $m = 51200$ (empirically determined on a holdout document collection that we did not include in this article). The dimensionality of the document representations $k_d \in \{64, 128, 256\}$ and the n-gram size $n \in \{4, 6, 8, 10, 12, 16, 24, 32\}$ are optimized on the validation set. Similar to d2v, models are trained for 15 iterations on the training set and we select the model iteration that performs best on the validation set; a single iteration consists of $\lceil \frac{1}{m} \sum_{d \in D} (|d| - n + 1) \rceil$ batches.

In addition, we consider lexical language models (QLM) (Zhai and Lafferty 2004) using the Indri engine with both Dirichlet (d) and Jelinek-Mercer (jm) smoothing; smoothing hyperparameters $\mu \in \{125, 250, 500, 750, 1000, 2000, 3000, 4000, 5000\}$ and $\lambda \in \{x \mid k \in \mathbb{N}_{>0}, k \leq 20, x = k/20\}$, respectively, are optimized on the validation set. The retrieval effectiveness of QLM is provided as a point of reference in **RQ1**. For **RQ2**, the QLM is used as a lexical retrieval model that is fused with latent vector space models to provide a mixture of lexical and semantic matching.

For the latent vector spaces (d2v, LSI, LDA, LSE, and NVSM), the vocabulary size is limited to the top-60k most frequent words as per Van Gysel et al. (2016a), given that latent methods rely on word co-occurrence to learn latent relations. For doc2vec, word2vec, LSI, and LDA, we use the Gensim⁶ implementation; the neural representation learning methods use our open source CUDA implementation described in Section 3.3; again, see footnote 1.

4.3.2 Combinations of QLM and Latent Features. We combine individual rankers by performing a grid search on the weights of a linear combination using 20-fold cross validation on the test sets (Table 1). For QLM, feature values are the log-probabilities of the query given the document, while for the latent features (NVSM and w2v-si), we use the cosine similarity between query/document representations. For every feature weight, we sweep between 0.0 and 1.0 with increments of 0.0125 on the fold training set. Individual features are normalized per query such that their values lie between 0 and 1. We select the weight configuration that achieves highest Mean Average Precision on the training set and use that configuration to score the test set. During scoring of the fold test set, we take the pool of the top-1k documents ranked by the individual features as candidate set.

4.4 Evaluation Measures and Statistical Significance

To address **RQ1**, **RQ2**, and **RQ3**, we report Mean Average Precision at rank 1000 (MAP@1000⁷), Normalized Discounted Cumulative Gain at rank 100 (NDCG@100), and Precision at rank 10 (P@10) to measure retrieval effectiveness. For **RQ3**, we also look at the per-query `titlestat_rel` (Buckley et al. 2007), the expected normalized term overlap between query and document. All evaluation measures are computed using TREC’s official evaluation tool, `trec_eval`.⁸ Wherever reported, significance of observed differences is determined using a two-tailed paired Student’s t-test (Smucker et al. 2007) (** $p < 0.01$; * $p < 0.05$; $p < 0.1$). For correlation coefficients, significance is determined using a permutation test ($\dagger p < 0.01$). For **RQ4**, we use Welch’s t-test to determine whether the mean L2-norm of mid-frequency (middle-50%) words is significantly different from the mean L2-norm of low- (bottom 25%) and high-frequency (top 25%) words.

5 RESULTS

First, we present a comparison between methods (**RQ1**) on ad hoc document retrieval, followed by a per-query comparison between methods (**RQ2**) and a combination experiment where we combine latent features with the lexical QLM (**RQ3**). We then relate regularities learned by the model to traditional retrieval statistics (**RQ4**).

⁶<https://github.com/RaRe-Technologies/gensim>.

⁷MAP@1000 summarizes the precision/recall curve and thus represents the tradeoff between precision and recall at different rank cutoffs. When using latent vector space models for generating a set of candidate documents (e.g., for subsequent re-ranking), the precision/recall tradeoff translates into precision for a particular candidate set size. A smaller candidate set is preferred as re-ranking many candidate documents can be expensive.

⁸https://github.com/usnistgov/trec_eval.

Table 2. Comparison of NVSM with Lexical (QLM with Dirichlet and Jelinek-Mercer Smoothing) and Latent (d2v, w2v, LSI, LDA, and LSE) Retrieval Models (Section 4.3) on Article Search Benchmarks (Section 4.2)

	AP88-89			FT			LA		
	MAP	NDCG	P@10	MAP	NDCG	P@10	MAP	NDCG	P@10
Bag-of-words features (for reference)									
QLM (jm)	0.199	0.346	0.365	0.218	0.356	0.283	0.182	0.331	0.221
QLM (d)	0.216	0.370	0.392	0.240	0.381	0.296	0.198	0.348	0.239
Latent features (for comparison)									
d2v	0.002	0.011	0.009	0.001	0.004	0.004	0.006	0.020	0.014
LDA	0.039	0.077	0.078	0.009	0.028	0.013	0.004	0.015	0.010
LSI	0.131	0.228	0.242	0.029	0.068	0.052	0.033	0.079	0.061
LSE	0.144	0.281	0.299	0.060	0.140	0.123	0.067	0.152	0.105
w2v-add	0.216	0.370	0.393	0.125	0.230	0.195	0.105	0.212	0.159
w2v-si	0.230	0.383	0.418	0.141	0.250	0.204	0.131	0.242	0.179
NVSM	0.257**	0.418**	0.425	0.172**	0.302***	0.239*	0.166**	0.300***	0.209*
	NY			Robust04			WSJ		
	MAP	NDCG	P@10	MAP	NDCG	P@10	MAP	NDCG	P@10
Bag-of-words features (for reference)									
QLM (jm)	0.158	0.270	0.376	0.201	0.359	0.369	0.175	0.315	0.345
QLM (d)	0.188	0.318	0.486	0.224	0.388	0.415	0.204	0.351	0.398
Latent features (for comparison)									
d2v	0.081	0.179	0.272	0.068	0.177	0.194	0.003	0.015	0.011
LDA	0.009	0.027	0.022	0.003	0.010	0.009	0.038	0.082	0.076
LSI	0.020	0.041	0.040	0.022	0.059	0.060	0.101	0.181	0.207
LSE	0.000	0.000	0.000	0.013	0.050	0.054	0.098	0.208	0.245
w2v-add	0.081	0.160	0.216	0.075	0.177	0.194	0.175	0.322	0.372
w2v-si	0.092	0.173	0.220	0.093	0.208	0.234	0.185	0.330	0.391
NVSM	0.117	0.208	0.296*	0.150***	0.287***	0.298***	0.208**	0.351	0.370

Significance (Section 4.4) is computed between w2v-si and NVSM. Bold values indicate the highest measure value for latent features.

5.1 Performance of NVSM

RQ1 Table 2 shows the retrieval results for ad hoc document retrieval on the newswire article collections (Section 4.2).

We see that NVSM outperforms all other latent rankers on all benchmarks. In particular, NVSM significantly outperforms (MAP@1000) the word2vec-based method that weighs word vectors according to self-information (significance is not achieved on NY). This is an interesting observation as NVSM is trained from scratch without the use of hand-engineered features (i.e., self-information). In the case of LSE, we see that its retrieval effectiveness diminishes. This is due to the large size of the NY collection (Table 1) and the lack of term specificity within LSE (§5.4).

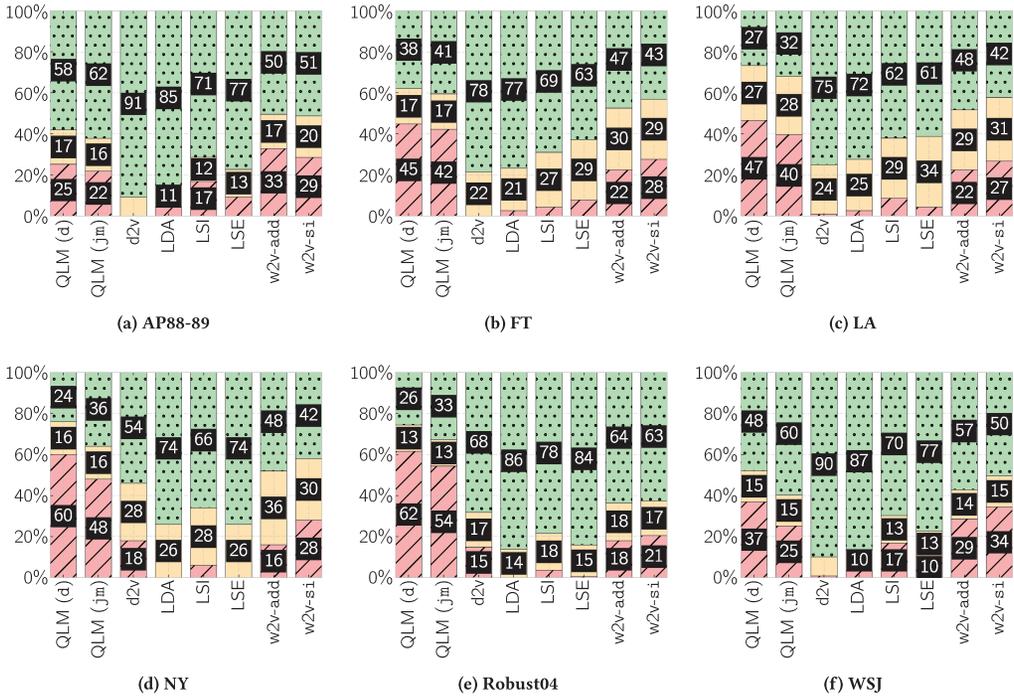


Fig. 1. Per-query pairwise ranker comparison between NVSM and the QLM (d), QLM (jm), d2v, w2v, LSI, LDA, and LSE rankers. For every bar, the dotted/green area, solid/orange, and red/slashed areas, respectively, depict the portion of queries for which NVSM outperforms, ties, or loses against the other ranker. One ranker outperforms the other if the absolute difference in MAP@1000 between both rankers exceeds δ .

Compared to the lexical QLM, NVSM performs better on the AP88-89 and WSJ benchmarks. However, it is known that no single ranker performs best on all test sets (Liu 2011; Shaw et al. 1994). In addition, NVSM is a latent model that performs a different type of matching than lexical models. Therefore, we first examine the per-query differences between rankers (RQ2) and later we will examine the complementary nature of the two types of matching by evaluating combinations of different ranking features (RQ3).

5.2 Query-Level Analysis

RQ2 Figure 1 shows the distribution of queries where one individual ranker performs better than the other ($|\Delta\text{MAP}@1000| > \delta$).

We observe similar trends across all benchmarks where NVSM performs best compared to all latent rankers. One competing vector space model, w2v-si, stands out as it is the strongest baseline and performs better than NVSM on 20–35% of queries, however, NVSM still beats w2v-si overall and specifically on 40–55% of queries. Moreover, Figure 1 shows us that QLM and NVSM make very different errors. This implies that the combination of QLM, w2v-si, and NVSM might improve performance even further.

While answering RQ1, we saw that for some benchmarks, latent methods (i.e., NVSM) perform better than lexical methods. While the amount of semantic matching needed depends on various factors, such as the query intent being informational or navigational/transactional (Broder 2002), we do see in Figure 1 that NVSM performs considerably better among latent methods in cases

where latent methods perform poorly (e.g., Robust04 in Figure 1(e)). Can we shine more light on the difference between NVSM and existing latent methods? We answer this question in the second half of the next section.

5.3 Semantic vs. Lexical Matching

RQ3 Beyond individual rankers, we now also consider combinations of the two best-performing vector space models with QLM (Shaw et al. 1994) (see Section 4.3.2 for details) in Table 3.

If we consider the QLM paired with either w2v-si or NVSM, we see that the combination involving NVSM outperforms the combination with w2v-si on four out of six benchmarks (AP88-89, LA, NY, Robust04). However, Figure 1 shows that NVSM and w2v-si outperform each other on different queries as well. Can we use this difference to our advantage?

The addition of NVSM to the QLM + w2v-si combination yields an improvement in terms of MAP@1000 on all benchmarks. Significance is achieved in five out of six benchmarks. In the case of NY and Robust04, the combination of all three rankers (QLM + w2v-si + NVSM) performs at about the same level as the combination of QLM + NVSM. However, the addition of NVSM to the QLM + w2v-si combination still creates a significant improvement over just the combination involving QLM and w2v only. For FT, the only benchmark where no significance is achieved, we do see that the relative increase in performance nearly doubles from the addition of NVSM. Consequently, we can conclude that the NVSM adds an additional matching signal.

Let us return to the question raised at the end of the previous section (Section 5.2): What exactly does the NVSM add in terms of content matching? We investigate this question by determining the amount of semantic matching needed. For each query, we compute `titlestat_rel` (Section 4.4), the expected normalized overlap between query terms and the terms of relevant document. If `titlestat_rel` is close to 1.0 for a particular query, then the query requires mostly lexical matching; on the other hand, if `titlestat_rel` is near 0.0 for a query, then none of the query's relevant document contain the query terms and semantic matching is needed. We continue by examining the per-query pairwise difference ($\Delta_{\text{MAP@1000}}$) between NVSM and the remaining lexical (QLM) and latent (d2v, LDA, LSI, LSE) features. If $\Delta_{\text{MAP@1000}} > 0$, then NVSM performs better than the other method and vice versa if $\Delta_{\text{MAP@1000}} < 0$. Table 4 shows the Pearson correlation between `titlestat_rel` and $\Delta_{\text{MAP@1000}}$. A positive correlation, as is the case for d2v, w2v, LDA, LSI, and LSE, indicates that NVSM performs better on queries that require lexical matching. Conversely, a negative correlation, such as observed for both variants of QLM, indicates that QLM performs better on queries that require lexical matching than NVSM. Combining this observation with the conclusion to **RQ2** (i.e., NVSM generally improves upon latent methods), we conclude that, in addition to semantic matching, NVSM also performs well in cases where lexical matching is needed and thus contributes a hybrid matching signal.

5.4 NVSM and Luhn Significance

If NVSM performs better at lexical matching than other latent vector space models, does it then also contain regularities associated with term specificity?

RQ4 Figure 2 shows the L2-norm of individual term representations for LSE (left scatter plots) and NVSM (right scatter plots).

Luhn (1958) measures the significance of words based on their frequency. They specify a lower and upper frequency cutoff to exclude frequent and infrequent words. For NVSM (scatter plot on the right for every benchmark), we find that infrequent and frequent terms have a statistically signif-

Table 3. Evaluation of Latent Features (d2v, w2v, LSI, LDA, LSE, and NVSM) as a Complementary Signal to QLM (d) (Section 4.3.2)

	AP88-89		
	MAP	NDCG	P@10
QLM (d)	0.216	0.370	0.392
QLM (d) + w2v-si	0.279 (+29%)	0.437 (+18%)	0.450 (+14%)
QLM (d) + NVSM	0.289 (+33%)	0.444 (+20%)	0.473 (+20%)
QLM (d) + w2v-si + NVSM	0.307*** (+42%)	0.466*** (+26%)	0.498*** (+27%)
	FT		
	MAP	NDCG	P@10
QLM (d)	0.240	0.381	0.296
QLM (d) + w2v-si	0.251 (+4%)	0.393 (+3%)	0.313 (+6%)
QLM (d) + NVSM	0.251 (+4%)	0.401 (+5%)	0.322 (+9%)
QLM (d) + w2v-si + NVSM	0.258 (+7%)	0.406 (+6%)	0.322 (+9%)
	LA		
	MAP	NDCG	P@10
QLM (d)	0.198	0.348	0.239
QLM (d) + w2v-si	0.212 (+7%)	0.360 (+3%)	0.236 (-1%)
QLM (d) + NVSM	0.220 (+11%)	0.376 (+7%)	0.244 (+1%)
QLM (d) + w2v-si + NVSM	0.226*** (+14%)	0.378*** (+8%)	0.250** (+4%)
	NY		
	MAP	NDCG	P@10
QLM (d)	0.188	0.318	0.486
QLM (d) + w2v-si	0.206 (+9%)	0.333 (+4%)	0.494 (+1%)
QLM (d) + NVSM	0.222 (+18%)	0.355** (+11%)	0.520 (+6%)
QLM (d) + w2v-si + NVSM	0.222*** (+18%)	0.353 (+10%)	0.526** (+8%)
	Robust04		
	MAP	NDCG	P@10
QLM (d)	0.224	0.388	0.415
QLM (d) + w2v-si	0.232 (+3%)	0.399 (+2%)	0.428 (+2%)
QLM (d) + NVSM	0.247 (+10%)	0.411 (+6%)	0.448*** (+7%)
QLM (d) + w2v-si + NVSM	0.247*** (+10%)	0.412*** (+6%)	0.446 (+7%)
	WSJ		
	MAP	NDCG	P@10
QLM (d)	0.204	0.351	0.398
QLM (d) + w2v-si	0.254 (+24%)	0.410 (+16%)	0.454 (+13%)
QLM (d) + NVSM	0.248 (+21%)	0.396 (+12%)	0.425 (+6%)
QLM (d) + w2v-si + NVSM	0.271*** (+32%)	0.426*** (+21%)	0.456 (+14%)

Significance (Section 4.4) is computed between QLM (d) + w2v-si and the best performing combination.

icant ($p < 0.01$), smaller L2-norm than terms of medium frequency (Section 4.4). This observation is further motivated by the shape of the relation between collection frequency in the collection and the L2-norm of term representations in Figure 2. The key observation that—within NVSM representations—terms of medium frequency are of greater importance (i.e., higher L2-norm) than low- or high-frequency terms closely corresponds to the theory of Luhn significance. Particularly noteworthy is the fact that the NVSM learned this relationship from an unsupervised objective directly, without any notion of relevance. The scatter plots on the left for every benchmark in

Table 4. Correlation Coefficients between `titlestat_rel` and $\Delta_{\text{MAP}@1000}$ between NVSM and the Other Methods

	AP88-89	FT	LA	NY	Robust04	WSJ
Bag-of-words features						
QLM (jm)	-0.102	-0.355 [†]	-0.188 [†]	-0.375 [†]	-0.337 [†]	-0.168 [†]
QLM (d)	-0.211 [†]	-0.456 [†]	-0.206 [†]	-0.432 [†]	-0.374 [†]	-0.275 [†]
Latent features						
d2v	0.625 [†]	0.415 [†]	0.473 [†]	0.195	0.161 [†]	0.459 [†]
LDA	0.545 [†]	0.406 [†]	0.497 [†]	0.380 [†]	0.309 [†]	0.376 [†]
LSI	0.420 [†]	0.376 [†]	0.528 [†]	0.344 [†]	0.217 [†]	0.282 [†]
LSE	0.506 [†]	0.361 [†]	0.396 [†]	0.374 [†]	0.289 [†]	0.270 [†]
w2v-add	0.374 [†]	0.275 [†]	0.450 [†]	0.283 [†]	0.230 [†]	0.393 [†]
w2v-si	0.232 [†]	0.171 [†]	0.316 [†]	0.203	0.247 [†]	0.357 [†]

A positive correlation indicates that NVSM is better at lexical matching, while a negative correlation indicates that NVSM is worse at lexical matching than the alternative. Significance (Section 4.4) is computed using a permutation test.

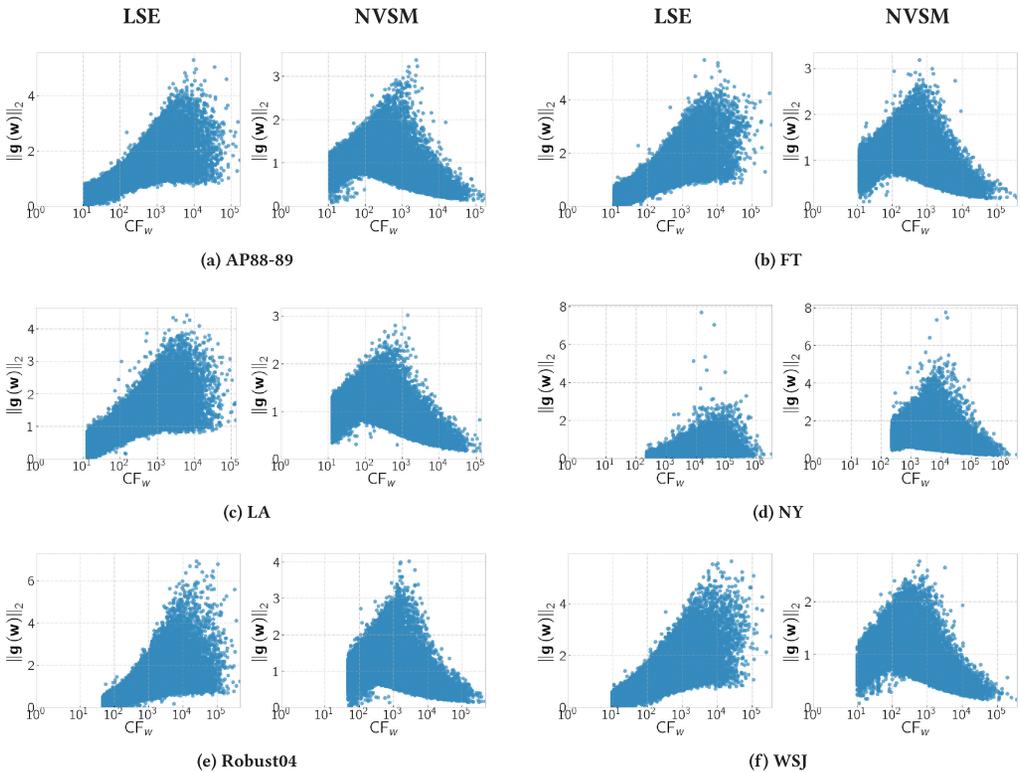


Fig. 2. Scatter plots of term frequency in the document collections and the L2-norm of the LSE (left) and NVSM (right) representations of these terms. In the case of LSE (left scatter plots for every benchmark), we observe that the L2-norm of term representations grows linearly with the term collection frequency, and, consequently, high-frequency terms are of greater importance within LSE representations. For NVSM (right scatter plot for every benchmark), we observe that terms of mid-frequency (middle 50%) have a statistically significant ($p < 0.01$) higher L2-norm (Section 4.4) and, consequently, are of greater importance for retrieval.

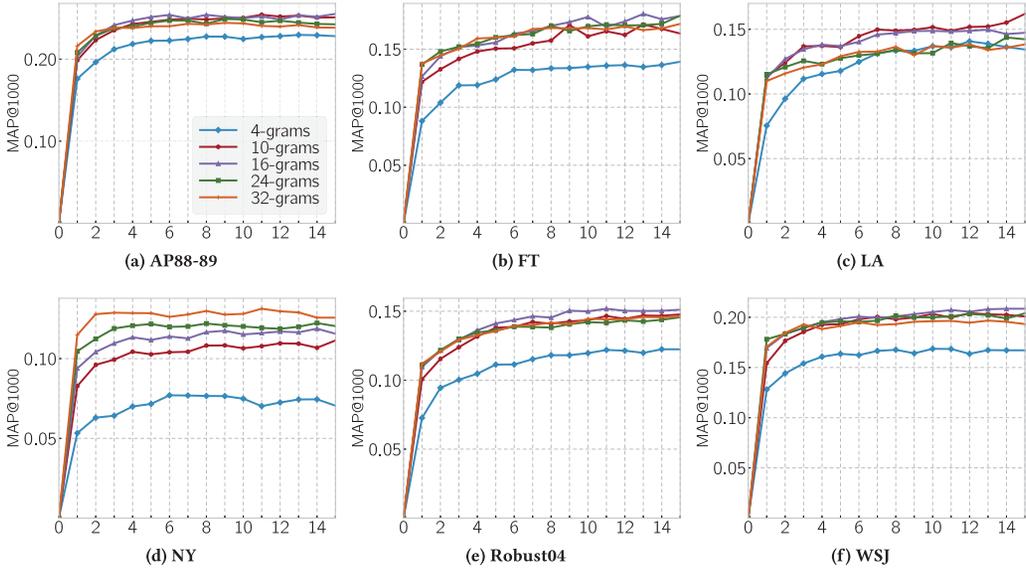


Fig. 3. Test set MAP@1000 as training progresses on article search benchmarks with document space dimensionality $k_d = 256$. We see that MAP@1000 converges to a fixed performance level with differently sized n -grams (here we show $n = 4, 10, 16, 24, 32$; the curves for the remaining values $n = 6, 8, 12$ are qualitatively similar and omitted to avoid clutter).

Figure 2 shows the same analysis for LSE term representations. Unlike with NVSM, we observe that the L2-norm of term representations grows linearly with the term collection frequency, and consequently, high-frequency terms are of greater importance within LSE representations. Therefore, the key difference between NVSM and LSE is that NVSM learns to better encode term specificity.

6 UNSUPERVISED DEPLOYMENT

In our experiments, we use a validation set for model selection (training iterations and hyperparameters). However, in many cases relevance labels are unavailable. Fortunately, NVSM learns representations of the document collection directly and does not require query-document relevance information. How can we choose values for the hyperparameters of NVSM in the absence of a validation set? For the majority of hyperparameters (Section 4.3) we follow the setup of previous work (Van Gysel et al. 2016a). We are, however, still tasked with the problem of choosing:

- (1) The number of training iterations
- (2) The dimensionality of the document representations k_d
- (3) The size of the n -grams used for training

We choose the dimensionality of the document representations $k_d = 256$ as the value was reported to work well for LSE (Van Gysel et al. 2016a). Figure 3 shows that MAP@1000 converges as the number of training iterations increases for different n -gram widths. Therefore, we train NVSM for 15 iterations and select the last iteration model.

The final remaining question is the choice of n -gram size used during training. This parameter has a big influence on model performance as it determines the amount of context from which semantic relationships are learned. Therefore, we propose to combine different vector spaces trained using different n -gram widths as follows. We write \mathcal{N} for the set of all k for which we construct

Table 5. Comparison with Single Cross-Validated NVSM and Ensemble of NVSM through the Unsupervised Combination of Models Trained on Differently Sized n -grams ($\mathcal{N} = \{2, 4, 8, 10, 12, 16, 24, 32\}$)

	AP88-89			FT		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.257	0.416	0.424	0.170	0.299	0.236
8 NVSMs (ensemble)	0.282***	0.453***	0.466***	0.212***	0.352***	0.282***
	LA			NY		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.168	0.302	0.211	0.117	0.208	0.296
8 NVSMs (ensemble)	0.188***	0.331***	0.230***	0.121	0.238***	0.324*
	Robust04			WSJ		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.150	0.287	0.297	0.208	0.351	0.372
8 NVSMs (ensemble)	0.171***	0.323***	0.331***	0.225***	0.385***	0.423***

Significance (Section 4.4) is computed between NVSM and the ensemble of NVSM.

an NVSM using k -grams. For a given query q , we rank documents $d \in D$ in descending order of

$$\text{score}_{\text{ensemble}}(q, d) = \sum_{k \in \mathcal{N}} \frac{\text{score}_{k\text{-grams}}(q, d) - \mu_{k\text{-grams}, q}}{\sigma_{k\text{-grams}, q}}, \quad (10)$$

where $\text{score}_{k\text{-grams}}(q, d)$ is Equation (3) for NVSM of k -grams and

$$\begin{aligned} \mu_{k\text{-grams}, q} &= \hat{\mathbb{E}} \left[\text{score}_{k\text{-grams}}(q, D) \right], \\ \sigma_{k\text{-grams}, q} &= \sqrt{\hat{\mathbb{V}} \left[\text{score}_{k\text{-grams}}(q, D) \right]}, \end{aligned} \quad (11)$$

denote the sample expectation and sample variance over documents D that are estimated on the top-1000 documents returned by the individual models, respectively. That is, we rank documents according to the sum of the standardized scores of vector space models trained with different n -gram widths. The score aggregation in Equation (10) is performed without any *a priori* knowledge about the n -gram sizes. Table 5 lists the performance of the unsupervised ensemble, where every model was trained for 15 iterations, against a single cross-validated model. We see that the unsupervised ensemble always outperforms (significantly in terms of MAP@1000 for all benchmarks except NY) the singleton model. Hence, we can easily deploy NVSM without any supervision and, surprisingly, it will perform better than individual models optimized on a validation set.

7 CONCLUSIONS

We have proposed the NVSM, which learns representations of a document collection in an unsupervised manner. We have shown that NVSM performs better than existing latent vector space/bag-of-words approaches. NVSM performs lexical and semantic matching in a latent space. NVSM provides a complementary signal to lexical language models. In addition, we have shown that NVSM automatically learns a notion of term specificity. Finally, we have given advice on how to select values for the hyperparameters of NVSM. Interestingly, an unsupervised ensemble of multiple models trained with different hyperparameters performs better than a single cross-validated model.

The evidence that NVSM provides a notion of lexical matching tells us that latent vector space models are not limited to only semantic matching. While the framework presented in this

article focuses on a single unsupervised objective, additional objectives (i.e., document/document or query/document similarity) can be incorporated to improve retrieval performance.

LSE (Van Gysel et al. 2016a) improved the learning time complexity of earlier entity retrieval models (Van Gysel et al. 2016b) such that they scale to $\sim 100k$ retrievable items (i.e., entities). However, as shown in Table 2, LSE performs poorly on article retrieval benchmarks. In this article, we extend LSE and learn vector spaces of $\sim 500k$ documents that perform better than existing latent vector spaces. As mentioned in the introduction, the main challenge for latent vector spaces is their limited scalability to large document collections due to space complexity. The observation that retrieval is not only impacted by the vector space representation of the relevant document, but also of the documents surrounding it, raises non-trivial questions regarding the distribution of document vectors over multiple machines. While there have been efforts towards distributed training of neural models, the application of distributed learning algorithms is left for future work. The unsupervised objective that learns from word sequences is limited by its inability to deal with very short documents. While this makes the unsupervised objective less applicable in domains such as web search, unsupervised bag-of-words approaches have the opposite problem of degrading performance when used to search over long documents. With respect to incremental indexing, there is currently no theoretically sound way to obtain representations for new documents that were added to the collection after the initial estimation of a NVSM. In the case of LDA or LSI, representations for new documents can be obtained by transforming bag-of-words vectors to the latent space. However, as the LDA/LSI transformation to the latent space is not updated after estimating the LDA/LSI model using the initial set of documents, this procedure can be catastrophic when topic drift occurs. For doc2vec, one way to obtain a representation for a previously unseen document is to keep all parameters fixed and train the representation of the new document using the standard training algorithm (Řehurek and Sojka 2010). This approach can also be used in the case of LSE or NVSM. However, there are no guarantees that the obtained representation will be of desirable quality. In addition, the same problem remains as with the bag-of-words methods. That is, the previously mentioned incremental updating mechanism is likely to fail when topic drift occurs.

We hope that our work and the insights resulting from it inspires others to further develop unsupervised neural retrieval models. Future work includes adding additional model expressiveness through depth or width, in-depth analysis of the various components on NVSM, analysis of the learned representations and their combination in various combination frameworks (Section 2.1), and engineering challenges regarding the scalability of the training procedure.

ACKNOWLEDGMENTS

We thank Adith Swaminathan, Alexey Borisov, Tom Kenter, Hosein Azarbyonad, Mostafa Dehghani, Nikos Voskarides and Adam Holmes and the anonymous reviewers for their helpful comments.

REFERENCES

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Derek Chris Olah, Mike Schuster, Jonathon Shlens, Benoi Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Zheng Xiaoqiang. 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. (2015). https://www.tensorflow.org/about/bib#large_scale_machine_learning_on_heterogeneous_distributed_systems.

- Qingyao Ai, Liu Yang, Jiafeng Guo, and W. Bruce Croft. 2016a. Analysis of the paragraph vector model for information retrieval. In *ICTIR*. ACM, 133–142.
- Qingyao Ai, Liu Yang, Jiafeng Guo, and W. Bruce Croft. 2016b. Improving language estimation with the paragraph vector model for ad-hoc retrieval. In *SIGIR*. ACM, 869–872.
- James Allan, Donna Harman, Evangelos Kanoulas, Dan Li, Christophe Van Gysel, and Ellen Voorhees. 2017. TREC 2017 common core track overview. In *TREC*.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL*. 238–247.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *J. Mach. Learn. Res.* 3 (2003), 1137–1155.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 3 (2003), 993–1022.
- Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016a. A context-aware time model for web search. In *SIGIR*. ACM, 205–214.
- Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016b. A neural click model for web search. In *WWW*. International World Wide Web Conferences Steering Committee, 531–541.
- Leonid Boytsov, David Novak, Yury Malkov, and Nyberg Eric. 2016. Off the beaten path: Let's replace term-based retrieval with k-NN search. In *CIKM*. 1099–1108.
- Andrei Broder. 2002. A taxonomy of web search. *SIGIR Forum* 36, 2 (2002), 3–10.
- Chris Buckley, Darrin Dimmick, Ian Soboroff, and Ellen Voorhees. 2007. Bias and the limits of pooling for large collections. *Inform. Retr.* 10, 6 (2007), 491–508.
- Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to rank using gradient descent. In *ICML*. ACM, 89–96.
- Olivier Chapelle and Ya Zhang. 2009. A dynamic bayesian network click model for web search ranking. In *WWW*. ACM, 1–10.
- Minmin Chen. 2017. Efficient vector representation for documents through corruption. In *ICLR*.
- Ronan Collobert, Jason Weston, Leon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* 12, Aug (2011), 2493–2537.
- Thomas M. Cover and Joy A. Thomas. 2012. *Elements of Information Theory*. John Wiley & Sons.
- Nick Craswell, W. Bruce Croft, Jiafeng Guo, Bhaskar Mitra, and Maarten de Rijke. 2016. Neu-IR: The SIGIR 2016 workshop on neural information retrieval. In *SIGIR*. ACM, 1245–1246.
- Scott C. Deerwester, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, and Richard A. Harshman. 1990. Indexing by latent semantic analysis. *J. Assoc. Inf. Sci. Technol.* 41, 6 (1990), 391–407.
- Li Deng, Xiaodong He, and Jianfeng Gao. 2013. Deep stacking networks for information retrieval. In *ICASSP*. 3153–3157.
- Susan T. Dumais. 1995. Latent semantic indexing (LSI): TREC-3 Report. In *TREC*. NIST, 219–230.
- Mike Folk, Gerd Heber, Quincey Koziol, Elena Pourmal, and Dana Robinson. 2011. An overview of the HDF5 technology suite and its applications. In *EDBT/ICDT Workshop on Array Databases*. ACM, 36–47.
- Debasis Ganguly, Dwaipayan Roy, Mandar Mitra, and Gareth J. F. Jones. 2015. Word embedding based generalized language model for information retrieval. In *SIGIR*. ACM, 795–798.
- Vincent Garcia, Eric Debreuve, and Michel Barlaud. 2008. Fast k nearest neighbor search using GPU. In *CVPRW*. IEEE, 1–6.
- Alex Graves and Navdeep Jaitly. 2014. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*. 1764–1772.
- Caglar Gulcehre, Marcin Moczulski, Misha Denil, and Yoshua Bengio. 2016. Noisy activation functions. *arXiv:1603.00391* (2016).
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016a. A deep relevance matching model for ad-hoc retrieval. In *CIKM*. ACM, 55–64.
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016b. Semantic matching by non-linear word transportation for information retrieval. In *CIKM*. ACM, 701–710.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*. 297–304.
- Donna Harman. 1992. The darpa tipster project. *SIGIR Forum* 26, 2 (Oct. 1992), 26–28.
- Donna Harman. 1993. Document detection data preparation. In *TIPSTER TEXT PROGRAM: PHASE I: Proceedings of a Workshop held at Fredricksburg, Virginia, September 19-23, 1993*. ACL, 17–31.
- Donna Harman and Ellen Voorhees. 1996. Overview of the fifth text retrieval conference. In *TREC-5*. 500–238.
- Thomas Hofmann. 1999. Probabilistic latent semantic indexing. In *SIGIR*. ACM, 50–57.
- Po-sen Huang, N. Mathews Ave Urbana, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*. 2333–2338.

- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR abs/1502.03167* (2015). <http://arxiv.org/abs/1502.03167>.
- Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *SIGKDD*. ACM, 133–142.
- Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. 2015. An empirical exploration of recurrent network architectures. In *ICML*. 2342–2350.
- Tom Kenter, Alexey Borisov, and Maarten de Rijke. 2016. Siamese CBOW: Optimizing word embeddings for sentence representations. In *ACL*. 941–951.
- Tom Kenter, Alexey Borisov, Christophe Van Gysel, Mostafa Dehghani, Maarten de Rijke, and Bhaskar Mitra. 2017. Neural networks for information retrieval. In *SIGIR 2017*. ACM, 1403–1406.
- Tom Kenter and Maarten de Rijke. 2015. Short text similarity with word embeddings. In *CIKM*. ACM, 1411–1420.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR abs/1412.6980* (2014).
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*. 1097–1105.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *ICML*. 1188–1196.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *IEEE* 86, 11 (1998), 2278–2324.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *NIPS*. 2177–2185.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Trans. Assoc. Comput. Linguist* 3 (2015), 211–225.
- Hang Li and Jun Xu. 2014. Semantic matching in search. *Found. Trends in Inform. Retr.* 7, 5 (June 2014), 343–469.
- Tie-Yan Liu. 2011. *Learning to Rank for Information Retrieval*. Springer.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM J. R&D* 2 (1958), 159–165.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Distributed representations of words and phrases and their compositionality. In *NIPS*. 3111–3119.
- Tomas Mikolov, Greg Corrado, Kai Chen, and Jeffrey Dean. 2013b. Efficient estimation of word representations in vector space. arXiv 1301.3781. (2013), 12 pages.
- Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to match using local and distributed representations of text for web search. In *WWW*.
- Gordon E. Moore. 1998. Cramming more components onto integrated circuits. *Proc. IEEE* 86, 1 (1998), 82–85.
- Marius Muja and David G. Lowe. 2014. Scalable nearest neighbor algorithms for high dimensional data. *Pattern Anal. Mach. Intell.* 36, 11 (2014), 2227–2240.
- Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. 2016. Improving document ranking with dual word embeddings. In *WWW*. International World Wide Web Conferences Steering Committee, 83–84.
- Kezban Dilek Onal, Ye Zhang, Ismail Sengor Altıngövdü, Md Mustafizur Rahman, Pinar Karagoz, Alex Braylan, Brandon Dang, Heng-Lu Chang, Henna Kim, Quinten McNamara, Aaron Angert, Edward Banner, Vivek Khetan, Tyler McDonnell, An Thanh Nguyen, Dan Xu, Byron C. Wallace, Maarten de Rijke, and Matthew Lease. 2018. Neural information retrieval: At the end of the early years. *Inform. Retr. J.* (2018). To appear.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *EMNLP*. 1532–1543.
- Radim Řehurek and Petr Sojka. 2010. Software framework for topic modelling with large corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. ELRA, Valletta, Malta, 45–50. <http://is.muni.cz/publication/884893/en>.
- Stephen Robertson. 2004. Understanding inverse document frequency: On theoretical arguments for IDF. *J. Doc.* 60, 5 (2004), 503–520.
- Stephen E. Robertson and Steve Walker. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR*. 232–241.
- Hasim Sak, Andrew W. Senior, and Françoise Beaufays. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Interspeech*. 338–342.
- Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *Int. J. Approx. Reason.* 50, 7 (2009), 969–978.
- Joseph A. Shaw, Edward A. Fox, Joseph A. Shaw, and Edward A. Fox. 1994. Combination of multiple searches. In *TREC*. 243–252.
- Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. A latent semantic model with convolutional-pooling structure for information retrieval. In *CIKM*. 101–110.
- Mark D. Smucker, James Allan, and Ben Carterette. 2007. A comparison of statistical significance tests for information retrieval evaluation. In *CIKM*. ACM, 623–632.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *J. Doc.* 28, 1 (1972), 11–21.

- Trevor Strohman, Donald Metzler, Howard Turtle, and W. Bruce Croft. 2005. Indri: A language model-based search engine for complex queries. In *ICIA*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*. 3104–3112.
- Ke Tran, Arianna Bisazza, and Christof Monz. 2016. Recurrent memory network for language modeling. In *NAACL-HLT*. ACL, 321–331.
- TREC. 1992–1999. TREC1-8 Adhoc Track. http://trec.nist.gov/data/qrels_eng.
- Xinhui Tu, Jimmy Xiangji Huang, Jing Luo, and Tingting He. 2016. Exploiting semantic coherence features for information retrieval. In *SIGIR*. ACM, 837–840.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: A simple and general method for semi-supervised learning. In *ACL*. Association for Computational Linguistics, 384–394.
- Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2016a. Learning latent vector spaces for product search. In *CIKM*. ACM, 165–174.
- Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2017a. Pyndri: A python interface to the indri search engine. In *ECIR*. Springer, 744–748.
- Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2017b. Semantic entity retrieval toolkit. In *Neu-IR 2017*.
- Christophe Van Gysel, Maarten de Rijke, and Marcel Worring. 2016b. Unsupervised, efficient and semantic expertise retrieval. In *WWW*. ACM, 1069–1079.
- Christophe Van Gysel, Bhaskar Mitra, Matteo Venanzi, Roy Rosemarin, Grzegorz Kukla, Piotr Grudzien, and Nicola Cancedda. 2017c. Reply with: Proactive recommendation of email attachments. In *CIKM*. ACM, 327–336.
- C. J. van Rijsbergen. 1979. *Information Retrieval* (2nd ed.). Butterworth-Heinemann.
- Michel Vidal-Naquet and Shimon Ullman. 2003. Object recognition with informative features and linear classification. In *ICCV*. IEEE, 281.
- Ellen M. Voorhees. 2005. The TREC robust retrieval track. *SIGIR Forum* 39, 1 (June 2005), 11–20.
- Ivan Vulić and Marie-Francine Moens. 2015. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In *SIGIR*. ACM, 363–372.
- Xing Wei and W. Bruce Croft. 2006. LDA-based document models for ad-hoc retrieval. In *SIGIR*. ACM, 178–185.
- Wikipedia. 2017. List of Nvidia graphics processing units — Wikipedia, The Free Encyclopedia. Retrieved August 8, 2017 from https://en.wikipedia.org/w/index.php?title=List_of_Nvidia_graphics_processing_units&oldid=792964538.
- Hamed Zamani and W. Bruce Croft. 2016a. Embedding-based query language models. In *ICTIR*. ACM, 147–156.
- Hamed Zamani and W. Bruce Croft. 2016b. Estimating embedding vectors for queries. In *ICTIR*. ACM, 123–132.
- Chengxiang Zhai and John Lafferty. 2004. A study of smoothing methods for language models applied to information retrieval. *ACM Trans. Inf. Syst.* 22, 2 (2004), 179–214.
- Guido Zuccon, Bevan Koopman, Peter Bruza, and Leif Azzopardi. 2015. Integrating and evaluating neural word embeddings in information retrieval. In *ADCS*. ACM, Article 12, 8 pages.

Received August 2017; revised January 2018; accepted March 2018