A Fast Deep Learning Model for Textual Relevance in Biomedical Information Retrieval

Sunil Mohan*

Chan Zuckerberg Initiative
Palo Alto, CA, USA
smohan@chanzuckerberg.com

Sun Kim

National Center for Biotechnology Information
Bethesda, MD, USA
sun.kim@nih.gov

Nicolas Fiorini

National Center for Biotechnology Information
Bethesda, MD, USA
nicolas.fiorini@nih.gov

Zhiyong Lu

National Center for Biotechnology Information
Bethesda, MD, USA
zhiyong.lu@nih.gov

ABSTRACT

Publications in the life sciences are characterized by a large technical vocabulary, with many lexical and semantic variations for expressing the same concept. Towards addressing the problem of relevance in biomedical literature search, we introduce a deep learning model for the relevance of a document's text to a keyword style query. Limited by a relatively small amount of training data, the model uses pre-trained word embeddings. With these, the model first computes a variable-length Delta matrix between the query and document, representing a difference between the two texts, which is then passed through a deep convolution stage followed by a deep feed-forward network to compute a relevance score. This results in a fast model suitable for use in an online search engine. The model is robust and outperforms comparable state-of-the-art deep learning approaches.

CCS CONCEPTS

• **Information systems** → **Learning to rank**; *Probabilistic retrieval models*; • **Computing methodologies** → **Learning to rank**; **Supervised learning by classification**; **Neural networks**; • **Applied computing** → **Life and medical sciences**;

KEYWORDS

Deep Learning; Biomedical Information Retrieval; Search; Learning to Rank

ACM Reference Format:

Sunil Mohan, Nicolas Fiorini, Sun Kim, and Zhiyong Lu. 2018. A Fast Deep Learning Model for Textual Relevance in Biomedical Information Retrieval. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3178876.3186049>

*Corresponding author. This work was done while this author was at NCBI.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186049>

1 INTRODUCTION

PubMed¹ is a free online search engine covering over 27 million articles from biomedical and life sciences journals and other texts, with about 1 million added each year. It is used worldwide by biomedical researchers, add healthcare professionals as well as lay people, serving about 3 million queries a day [5]. While expert users either search for most recent articles by an author or construct elaborate query expressions, most queries are short keyword-ese, covering one or two biomedical concepts. Although the size of the corpus is much smaller than in general web search, biomedical literature uses a very large technical vocabulary (e.g. the UMLS² metathesaurus [1] specifies over 3 million biomedical concepts, along with several lexical variations and synonymous phrases. This makes it much harder to identify concepts across documents (e.g. see [18]). To improve the retrieval, PubMed expands a user's query by mapping it to related MeSH[®] terms [22]. While this increases recall, it often decreases precision [12]. Usage analysis [5] shows that PubMed users are persistent, often reformulating their query, issuing over 4 queries per session on average. As part of improving relevance for such keyword queries, we describe a deep learning model that addresses the relevance of a document's text to the query. The eventual goal is for this model to be incorporated as a factor into a reranker that also includes other document attributes and metadata (e.g. year, journal).

To train our model, we collected data from PubMed click logs, restricting this to relevance search instead of the default sort order by date. Removing author searches and disjunctive boolean expressions resulted in a training set of about 20k queries. Given the small size of this data, we pre-trained word embeddings using word2vec [26] on the entire PubMed corpus, producing a vocabulary of about 200k. This large gap between training data and vocabulary sizes highlights a major challenge: how to make the model robust? Our Delta deep learning model begins by computing a variable-sized 'Delta' matrix between a document and a query, comprising the vector difference between document word embeddings and the closest matching query word, and three scalar similarity measures between the query and document. The document is truncated to control run-time cost. The Delta matrix is processed through a stacked convolutional network and pooled to a fixed length. This, together with a summary query match statistic, is processed by a

¹<http://pubmed.gov>

²<http://umlsk.nlm.nih.gov>

feed-forward network to produce a relevance score. Pairwise loss is optimized for training. This approach produces a model that is both robust, and fast enough for use in a search engine.

In addition to model robustness, we also wanted to address two common search engine problems: (i) the *under-specified query problem* [7], where even irrelevant documents have prominent presence of the query terms, and relevance requires analysis of the topics and semantics not directly specified in the query, and (ii) the *term mismatch problem* [8], which requires detection of related alternative terms or phrases in the document when the actual query terms are not in the document. Our experiments show the Delta model outperforms traditional lexical match factors and some related state-of-the-art neural approaches.

The next sections discuss some related work, followed by a description of the model, the experiments and evaluation of results, ending with some concluding remarks.

2 RELATED WORK

Traditional lexical Information Retrieval (IR) factors, like Okapi BM25 [33] and Query Likelihood [27], measure the prominence of query terms occurring in documents treated as bags of words. Neural approaches to text relevance attempt to go beyond exact matches of query terms in documents, and model a degree of semantic match as a complex function in a continuous space (good reviews can be found in [28, 40]). We will discuss some related approaches here.

Most neural models begin by mapping words to points embedded in a real space. A popular approach (e.g. [11, 13]), also used in our model, is to pre-train word embeddings, e.g. using word2vec [25, 26]. The benefit of this approach is that a much larger unlabeled corpus can be used to train the embeddings, and our ‘Delta matrix’ takes advantage of the semantic relationships captured in the vector differences between words.

The simplest embeddings based model is Word Mover’s Distance (WMD) [19], a non-parameterized model for text similarity that does not require any training. We use this as one of our baselines. Pre-trained embeddings are not necessarily targeted for optimal relevance scores. Nalisnick et al. [31] also use the ‘input’ vectors normally discarded by word2vec to overcome some of this limitation. The ‘DSSM’ models of [14, 36] take a different approach by mapping each word to a bag of letter tri-grams and combining the corresponding one-hot vectors. Xiong et al. [38] show that training word embeddings as part of the relevance model has a major impact on the performance of a relevance model. However this requires a large amount of training data. Diaz et al. [4] show that ‘locally trained’ embeddings on pseudo-relevance feedback documents can provide better results, while admitting that this approach is not “computationally convenient”.

Neural relevance models also differ in how they process document and query text. Some (e.g. [9, 13, 14, 36]) process each document and query using separate ‘Siamese’ networks into independent semantic vectors. A second stage then scores the similarity between these vectors. This approach is very attractive for search engines, because the document vectors can be pre-processed and stored, and the query vector need be produced once before scoring the

Figure 1: The Delta Relevance Model.

documents, significantly reducing the cost at query time. We use the recent model described in [35] as a baseline.

Another approach to text matching first develops ‘local interactions’ by comparing all possible combinations of words and word sequences between the document and query texts, often starting with a document-query word similarity matrix. Examples are described in [11, 13, 23, 32, 38]. The authors in [11] argue that the local interaction based approach is better at capturing detail, especially exact query term matches, and in their experiments their ‘DRMM’ model outperforms many previous approaches. This is a more computationally intensive architecture that does not allow any pre-computation. We take a similar approach by pairing each document word with a single query word, followed by deep convolutions to capture some related compositional semantics. Run-time cost in our approach is controlled by truncating the document. We show that our approach outperforms the DRMM model.

The ‘PACRR-firstk’ model in [15] also truncates the document, then processes the resulting similarity matrix through several 2D convolutional layers aimed at capturing n-gram similarities, followed by a recurrent layer, resulting in a fairly complex model. The ‘DUET’ model described in [29] combines a local interaction model with an independent semantic vector model, with the goal of combining the benefits of ‘exact match’ and embedding based word similarities. Our simpler approach explicitly targets run-time efficiency, and a variant of the Delta model combines some lexical factors (similar to [35]) to further improve ranking performance.

3 THE DELTA MODEL

The components of the Delta Relevance Model (figure 1) are described below. The unshaded blocks represent inputs to the model: two vectors of word indices, one each for the Document D and the Query Q , and a vector of query-document Lexical Match factors $L_{DQ} = \text{lexmatch}(D, Q)$ for some chosen lexical match function.

The small size of the training data ($\sim 20,000$ queries) compared to the vocabulary size ($\sim 200,000$) prevented us from training word embeddings as part of the model training. We had to adapt the word vectors pre-trained using word2vec’s unsupervised approach, to the task of relevance prediction. The Delta model uses two techniques that help in this and thus learn a richer and more robust decision surface. Changing the input space from word embeddings to differences in word embeddings shifts the domain of the decision surface to coordinates relative to the query. In addition, the Delta model’s use of a stack of convolution layers instead of a single layer adds more non-linearities to help capture a complex decision surface, a

technique successful in image recognition [37]. The convolution layers also extract relevance-match signals from text n -grams, and are much faster than a recurrent layer which has a similar goal.

3.1 Word Embeddings

We leveraged the large PubMed corpus of over 27 million documents to pre-train the word vectors, using the SkipGram Hierarchical Softmax method of word2vec [26], with a window size of ± 5 , a minimum term-frequency of 101, and a word-vector size of $V = 300$ (see [3] for experiments with different parameter settings for biomedical text). This resulted in a vocabulary of 207,716 words. Rare words were replaced with the generic “UNK” token, which was initialized to $\sim U[-0.25, 0.25]$, as in [35].

Given a document word sequence $D = \langle w_1^d, \dots, w_N^d \rangle$ and query text $Q = \langle w_1^q, \dots, w_M^q \rangle$, where w_i are indices into the vocabulary, the Embeddings layer replaces each word with its vector, giving us $D^e = \langle d_1, \dots, d_N \rangle, Q^e = \langle q_1, \dots, q_M \rangle$ where each $d_i, q_i \in \mathbb{R}^V$, and V is the size of the word embedding. If a document has fewer than N words, or the query fewer than M words, they are padded on the right with zeros. Longer documents are truncated, and M is the longest query length in our data (see section 4.1).

3.2 The Delta Stage

This is an unparameterized stage, responsible for computing the Delta Matrix between the Document and the Query as follows:

- (1) Compute the Euclidean distance between each pair d_i, q_j .
- (2) For each document word d_i , determine the closest query word q_i^* , using these distances.
- (3) Compute the vector differences $(d_i - q_i^*)$.
- (4) Compute the Delta features: cosine(d_i, q_i^*), $|d_i - q_i^*|$ and normalized proximity similarity metric $1 - |d_i - q_i^*|/(|d_i| + |q_i^*|)$.

The output of this stage is $\Delta = \langle \delta_1, \dots, \delta_N \rangle$, a $N \times (V + 3)$ sized real matrix. All operations above are masked to ignore padding.

3.3 The Network

The trainable portion of the Delta model consists of a Convolutional Stage followed by a Feed-Forward Stage. The Convolutional stage attempts to pick up significant contextual and n -gram similarity features. These are combined with the Lexical Match features L_{DQ} , then processed by the Feed-Forward Stage to produce a final relevance score. The weights for the layers in these stages comprise the trainable parameters for the model.

A convolution operation [10, 21] has the parameters: border mode, number of filters or feature maps n_f , filter width k and stride s . We use 1-dimensional convolution along the text width with border mode ‘same’ which implicitly pads the input on either side before convolving, and a stride $s = 1$, resulting in an output of the same width as the input. Using $\text{conv}(X; n_f, k)$ to represent such a convolution on input X , the Delta model’s Convolutional Stage performs the following operations:

$$Y^{(i)} = f_i(\text{conv}(Y^{(i-1)}; n_f, k)) \quad i = 1, \dots, N_C$$

where f_i is an activation function, $Y^{(0)} = \Delta$, N_C is the number of convolution layers, and $Y^{(N_C)} \in \mathbb{R}^{N \times n_f}$. The number of filters and

filter width are kept the same for each layer, and all operations are masked to ignore padding.

The output of the final convolution layer is ‘max-pooled’ by taking the maximum of each of the n_f output features along the text width dimension, yielding a vector of size n_f . This is the output of the Convolution Stage:

$$Y = \max(Y^{(N_C)}; \text{axis} = 0) \in \mathbb{R}^{n_f}$$

This output is combined with the Lexical Match features L_{DQ} and sent to the Feed-Forward stage which is a series of N_F layers:

$$Z^{(i)} = g_i(Z^{(i-1)} \cdot W_i + b_i) \quad i = 1, \dots, N_F$$

Here: g_i is an activation function, $Z^{(0)} = \langle Y, L_{DQ} \rangle$, $Z^{(i-1)} \cdot W_i$ is the matrix multiplication operation, N_F is the number of layers in the Feed-Forward stage, and W_i, b_i are sized so that the number of outputs of the i -th layer are K_i . The output of the final layer is a single number, i.e. $K_{N_F} = 1, Z^{(N_F)} \in \mathbb{R}$, representing the relevance score of the document D to the query Q .

3.4 Training The Model

The training data derived from PubMed’s click logs provides relevance levels for query-document pairs based on the number of clicks they received (see next section for more details). The Delta relevance model is trained to give more relevant documents a higher score by tuning its parameters Θ to minimize the pairwise maximum margin loss. Given a query Q and two matching documents D^+ , D^- where D^+ has higher relevance to the query than D^- , the loss for this triple is expressed as:

$$\mathcal{L}(Q, D^+, D^-; \Theta) = \max(0, 1 - s(Q, D^+; \Theta) + s(Q, D^-; \Theta))$$

where $s(Q, D; \Theta)$ is the relevance score produced by the Delta model for the query-document pair Q, D . The Adagrad [6] stochastic gradient descent method was used to train the model, using a mini-batch size of 256. In addition to early stopping, separate L2 regularization costs were added on the weights of the Convolutional and Feed-Forward stages, and a dropout layer was added before the max-pooling layer in the Convolutional stage. The regularization coefficients and dropout probability were tuned using the held out validation data. Adding dropout layers to the Feed-Forward stage was also tested but was not found to help.

4 EXPERIMENTAL SETUP

4.1 The Data

We collected query-document pairs extracted from PubMed click logs over several months where users selected ‘Best Match’ (relevance) as the retrieval sort order and clicked on at least one document in the search results. We recorded the first page of results of up to 20 documents, supplemented with the clicked document if it was not on the first page. Since our primary goal was to improve relevance for simple keyword style queries, we discarded queries containing disjunctive expressions, faceted queries, and queries longer than 7 words. Log extracts were further restricted to queries with at least 21 documents, and at least 3 clicked documents. These filters reduced the logs to about 33,500 queries, which were randomly split to 60% training, and 20% each for validation and testing.

	Full Test Data	Neg20+	OneNewWord	AllNewWords
Nbr. of Queries	6,734	2,600	1,732	933
Nbr. of Samples	413,971	208,723	86,438	47,825
Prop. of Samples +ive	45.2%	39.5%	49.2%	49.2%
Prop. of Samples -ive	54.8%	60.5%	50.8%	50.8%
+ives without all Query terms in Title	38.7%	13.9%	32.7%	22.9%
-ives with all Query terms in Title	59.5%	83.6%	68.2%	78.0%

Table 1: Test data and its subsets

Relevance Levels. The relevance level assigned to each query-document pair extracted from click logs is a probability of relevance, scaled to the range $[0, 100]$ so that the minimum possible non-zero relevance is 1. For each query-document pair Q, D , we accumulated over the collection period the number of click-throughs $c(D, Q)$ from search results to the document summary page, whether the document’s full-text was available in PubMed $I_{ft}(D) \in \{0, 1\}$, and the number of subsequent click-throughs $c_{ft}(D, Q)$ to the document’s full-text. These were used to derive a weighted click-count $c_w(Q, D)$ that rewarded documents for which full-text was requested without penalizing those for which full-text was not available. From that a probability of relevance $rel(Q, D)$ was calculated:

$$c_w(Q, D) = (\mu + (1 - I_{ft}(D)) \cdot \lambda) \cdot c(D, Q) + (1 - \mu) \cdot c_{ft}(D, Q)$$

$$rel(Q, D) = \frac{c_w(Q, D)}{\sum_{D'} c_w(Q, D')}$$

$$srel(Q, D) = 1 + 99 \times rel(Q, D) \dots \text{ if } rel(Q, D) > 0$$

$$= 0 \dots \text{ if } rel(Q, D) = 0$$

Finally, we scaled the non-zero relevance levels to the range $(1, 100]$ to get $srel(Q, D)$. This ensured a minimum margin between documents of low relevance and no relevance, and also put a high penalty in the NDCG metric for ranking high relevance documents below low relevance ones. The coefficients were tuned to match NCBI domain experts’ relevance judgments: $\mu = 0.333$, $\lambda = 0.067$.

Tokenization. Each document in our data had a Title and an Abstract. For the neural models, we concatenated these to form the document’s ‘Text’. All document and query text was tokenized by splitting on space and punctuation, while preserving abbreviations and numeric forms, followed by a conversion to lower-case. To further reduce the vocabulary size, all punctuation was discarded and numeric forms were collapsed into 7 classes: Integer, Fraction in $(0, 1)$, Real number, year “19xx”, year “20xx”, Percentage (number followed by “%”), and dollar amount (number preceded by “\$”). While word2vec processed the tokenized documents in sentences, the document input to the neural models was a flat sequence of words without sentence breaks or markers. The distribution of document text widths (nbr. of words) in the data is shown in figure 2a. We experimented with stopword removal in the Query and the Document, but they did not help.

Test Data Subsets. The 20% held out test data comprised 6,734 queries and 413,971 samples (query-document pairs). Presence of query words in a document’s Title is often a good indication of relevance. Among the relevant documents (“+ives”) for all the test queries, 38.7% did not contain all query terms in the title. Similarly

59.5% of all the non-relevant documents (“-ives”) actually contained all the query terms in their title (see table 1).

In addition to comparing ranking metrics of the different approaches on the test data, we wanted to explore model robustness, and model performance with *under-specified queries*. To help answer these questions, we also compared ranking metrics on the following subsets of the test data:

Neg20+: This consisted of all queries for which there were at least 20 non-relevant documents that contained all the query words in the title. This subset was used to evaluate performance on under-specified queries.

OneNewWord: The 1,732 test queries that contained at least one new word not occurring in any training or validation queries.

AllNewWords: A smaller subset of queries all of whose words were new: none of the training or validation queries included these words.

The last two subsets help evaluate model robustness. The statistics of the test data and its subsets are summarized in table 1.

4.2 Configuration Settings for the Delta Model

The Delta Model’s hyper-parameters were tuned to optimize the ranking metric NDCG₂₀ on the validation data. We found truncating documents to the first $N = 50$ words provided a good compromise between ranking performance and the run time to score a query-document pair (discussed below), with larger values providing only marginal improvements. The maximum query size was $M = 7$ as described above. The Convolutional stage used $N_C = 3$ layers of convolutions, each with a filter width $k = 3$. We report metrics for various number of filters n_f below. The Feed-Forward stage used $N_F = 3$ layers. Finally, we found downsampling the training data so that there were an equal number of relevant and non-relevant documents for each query to produce the best model, resulting in 7,084,244 training samples of (Q, D^+, D^-) triples. This downsampling was also performed for the other neural models described below. The validation and test data were not downsampled.

With the maximum-margin loss function, there was no reason to constrain the range of the final layer’s activation function. We got best results using the Leaky Rectified Linear Unit (Leaky ReLU) [24] with the slope of the negative region fixed at $\alpha = 0.3$. The Leaky ReLU was also used as the activation function for all the other layers of the Feed-Forward and Convolutional stages.

An earlier version of the Delta model is described in [30]. The main changes since then are: a simpler Delta Matrix, changes to the activation functions used in all the stages, training to a pairwise loss function with different sample weighting, and comprehensive test

on a number of lexical features. These changes resulted in a $\sim 10\%$ improvement in the NDCG metrics. We only report metrics for the current version of the Delta model below, along with a comparison against some new baselines.

4.2.1 Relevance-based Sample Weighting. Best results were obtained by adding a weight to each (Q, D^+, D^-) training sample in the loss function by taking the square-root of the difference in the scaled-relevance levels of the two documents:

$$\text{weight}(Q, D^+, D^-) = (\text{srel}(Q, D^+) - \text{srel}(Q, D^-))^{0.5}$$

4.2.2 Lexical Match Features. As an extension to the “word overlap measures” used in the SevMos model [35], we tested 18 features for use as the ‘Lexical Match Factors’ input to the Delta model:

- (1) Proportion of unique Query words present in document Text.
- (2) Proportion of unique Query bigrams present in doc Text.
- (3) Jaccard Similarity between Query and document Text.
- (4) IDF weighted version of (1).
- (5) An IDF weighted version of Jaccard Similarity (3).
- (6) BM25 on Query, document Title.
- (7) BM25 on Query, document Abstract.
- (8) BM25 on Query, document Text.
- (9) Proportion of unique Query words present in document Title.
- (10) Proportion of unique Query bigrams present in doc Title.
- (11) Jaccard Similarity between Query and document Title.
- (12) IDF weighted version of (9).
- (13) An IDF weighted version of Jaccard Similarity (11).
- (14) Proportion of unique Query words present in doc Abstract.
- (15) Proportion of unique Query bigrams present in doc Abstract.
- (16) Jaccard Similarity between Query and document Abstract.
- (17) IDF weighted version of (14).
- (18) An IDF weighted version of Jaccard Similarity (16).

To compute these factors, Queries and Documents were tokenized as described above, without the rare word conflation needed for computing word embeddings. Document Text refers to the combined Title and Abstract, each of these (as well as the Query) treated as a sequence of words with no truncation. These factors were selected based on the speed of their computation in a search engine. Factors (3, 5, 11, 13, 16, 18) were also used in [35].

4.3 Baselines

We compared the performance of the Delta deep learning model against some traditional bag-of-words based textual relevance factors, a distance measurement based on distributional representations of words, and a couple of recent neural network models.

4.3.1 Lexical Factors. We compared the performance of Okapi BM25 [33] on the document Title, Abstract and Text (Title + Abstract), and found BM25 on Title to give the best ranking performance, with parameter settings at $k_1 = 2.0$ and $b = 0.75$.

The second lexical factor we tested was Unigram Query Likelihood (UQLM), which estimates the probability with which the most likely random process that generated the bag-of-words representation of the document, would generate the query. It is based on a generative unigram language model that is a mixture of two multinomial models [27] based on the document and the corpus, combined using Dirichlet smoothing [20, 39]. Just like in the case

of BM25, we found UQLM applied to the document Title to perform the best, and quote only those metrics below.

4.3.2 Word Mover’s Distance. Since all the neural models in our experiments started with pre-trained word embeddings, the *Word Mover’s Distance* (WMD) model [19] for text dis-similarity (score decreases with increasing similarity) was an obvious baseline approach. Based on the *Earth Mover’s Distance* [34] applied to a bag-of-words representation for text, it is a non-parameterized approach to determine the minimum amount of total transportation cost (sum of product of inter-word cost and amount transported) needed to convert one document into the other. It uses the Euclidean distance between the word-vector representations of two words as the cost of moving from one word to another. We only report metrics for WMD applied to the document Title without removal of stop-words, as it performed better than the other alternatives tested.

4.3.3 The Severyn-Moschitti Model. As a recent example of the Independent Semantic Vector approach, we implemented the relevance classification model described in [35], along with a few variations. The query and document are fed into separate Convolutional stages, each comprising a single convolution layer with 256 feature maps and a filter width of 5, followed by Dropout and Global Max-Pooling. A similarity measure is computed from these pooled outputs using a similarity weight matrix. The similarity measure, the pooled output, and some lexical match features (“overlap measures” in [35]) are fed into a Classifier stage consisting of a series of feed-forward layers. In our experiments, we provided the SevMos models with all 18 lexical match features described in section 4.2.2. Optimal values for the L2-regularization and Dropout probability hyper-parameters were determined by tuning on validation data, as described for the Delta model.

We tested several variants of this model covering: replacing the single convolution layer with a 3-layer stack of convolutions of filter width 3, similar to the Delta model’s Convolutional stage; training the model as a classifier v/s a relevance scorer to the pairwise max-margin loss; and various sample-weighting schemes. Best results were obtained with the classification model using a 3-layer convolution stack and square-root weighting of samples. We report the metrics for this approach as the “SevMos-C3” model below, and the corresponding single convolution layer based classifier as the “SevMos-C1” model.

4.3.4 The DRMM Model. The Deep Relevance Matching Model (DRMM) is a recent example of the Local Interaction approach to text relevance, described in [11] to outperform several previous neural models on the Robust04 and ClueWeb-09-Cat-B datasets. While it is a simple model with only 162 trainable parameters, it begins by computing the cosine similarity between the embeddings of each document and query word pair, which dominates the model’s computational cost. We implemented the DRMM model as described in [11], using the Krovetz word stemmer during text tokenization, stopwords removed from queries, and the CBOW method of [26] to compute word embeddings.

We tested DRMM on increasing values of N (maximum document width) and found the ranking metrics stopped improving after a width of 200 words (figure 2b). DRMM uses the same pairwise loss function; we found different sample-weighting schemes to have

Figure 2: (a) Distribution of Document Text widths. (b) DRMM performance by max document width N .

an insignificant effect on the metrics. We report metrics for the version using square-root weighting and $N = 200$.

4.4 Metrics

Each of the following metrics has values in the range $[0, 1]$, with higher values for better rankings. Scoring ties in all the compared approaches were resolved by sorting on decreasing document-id.

4.4.1 NDCG. Discounted Cumulative Gain (DCG) [16] is a relevance and rank correlation metric that penalizes placement of relevant documents at lower ranks, computed as:

$$DCG(n) = \sum_{i=1}^n \frac{2^{rel(i)} - 1}{\log_2(i + 1)}$$

where n is the rank to which DCG is accumulated, and $rel(i) \geq 0$ is the relevance level of the document placed at rank i . Normalized Discounted Cumulative Gain (NDCG) then measures the relative DCG of a ranking compared to the best possible ranking for that data: $NDCG(n) = DCG(n)/IDCG(n)$, where $IDCG(n)$ is the DCG(n) for the ideal ranking. When there are multiple queries, NDCG refers to the mean value across queries. We use the scaled relevance levels (section 4.1), and quote “NDCG.20” metrics for $n = 20$.

	NDCG.20	MAP	Prec.5
rev DocID	0.141	0.455	0.344
BM25-Title ¹	0.325	0.567	0.591
UQLM-Title	0.314	0.560	0.574
WMD-Title ²	0.329 ⁼¹	0.579 ⁺¹	0.603 ⁺¹
DRMM	0.300 ⁻¹	0.545 ⁻¹	0.549 ⁻¹
SevMos-C1 ³	0.352 ⁺²	0.597 ⁺²	0.625 ⁺²
SevMos-C3 ⁴	0.373 ⁺³	0.594 ^{=3,+2}	0.626 ^{=3,+2}
Delta-32 ⁵	0.365 ^{+3,-4}	0.601 ^{+3,4}	0.634 ^{+3,4}
Delta-32-Lex3	0.394 ^{+4,5}	0.609 ^{+4,5}	0.646 ^{+4,5}

Table 2: Ranking metrics on the Full test data. The superscripts indicate statistical comparisons: ‘+’ for increase, ‘-’ for decrease, ‘=’ for equivalent, to a 99% confidence using a paired t-test. The comparison baselines are indicated with numbers 1 through 5, as marked in the first column. Highest values are in bold.

4.4.2 Precision at Rank and MAP. Average Precision [2] measures, for a single query, the precision observed in a ranking up to the rank of each relevant document, averaged over the number of relevant documents for that query. It is thus a ranking measure that factors out the size of the ranked list and the number of relevant documents, without any rank-based penalization or discounting. We quote the Mean Average Precision (MAP), which is the mean of the Average Precision across queries in our test dataset. We also quote some Precision at rank n metrics (“Prec. n ” in the tables).

5 EVALUATION

5.1 Test Metrics

We compare ranking performance of two versions of the Delta model against the other approaches. The ‘Delta-32’ model uses $n_f = 32$ feature maps in the Convolutional stage, and no Lexical Match features. The ‘Delta-32-Lex3’ version of the model adds the following three Lexical Match features: BM25 on the Document Abstract, IDF weighted Jaccard Similarity between the Query and the Document Title, and IDF-weighted proportion of unique Query words in the Document Title. These features were selected using greedy search on the list of 18 described earlier, with NDCG.20 on the validation data as the selection criterion. The feature selection was limited to three to control model run-time cost.

We begin with ranking metrics for the various approaches on the full test data (table 2). The first row in the table provides metrics for an uninformed ranker, where documents are ranked on decreasing document id, to provide a low threshold of performance. The table indicates whether there was a statistically significant change (to a 99% confidence, using a paired t-test) against a baseline. Among the non-trained relevance models, Word Mover’s Distance (WMD) performed at least as well as BM25, with no change in NDCG.20, but an improvement in MAP and Prec.5 (Precision at rank 5), while the Query Language Model (UQLM) did not match BM25’s level of performance. Among the trained neural models, DRMM performed the worst, with lower metrics than even BM25. The SevMos-C1 model performed better overall than WMD, and SevMos-C3 further improved the NDCG.20 score.

	NDCG.20	MAP	Prec.5
rev DocID	0.081	0.413	0.310
BM25-Title ¹	0.233	0.474	0.490
WMD-Title ²	0.243 ⁺¹	0.483 ⁺¹	0.496 ⁼¹
DRMM	0.242 ^{+1,=2}	0.461 ^{-1,2}	0.462 ^{-1,2}
SevMos-C1 ³	0.290 ⁺²	0.510 ⁺²	0.538 ⁺²
SevMos-C3 ⁴	0.304 ^{+2,3}	0.502 ^{+2,-3}	0.535 ^{+2,=3}
Delta-32 ⁵	0.296 ^{+2,=3}	0.513 ^{+2,=3}	0.550 ^{+2,3}
Delta-32-Lex3	0.326^{+4,5}	0.522^{+4,5}	0.560^{+4,5}

Table 3: Ranking metrics on the ‘Neg20+’ test data

	NDCG.20	MAP	Prec.5
rev DocID	0.191	0.488	0.364
BM25-Title ¹	0.333	0.593	0.604
WMD-Title ²	0.330 ⁼¹	0.603 ⁺¹	0.614 ⁺¹
DRMM	0.318 ⁻¹	0.580 ⁻¹	0.578 ⁻¹
SevMos-C1 ³	0.358 ⁺²	0.624 ⁺²	0.642 ⁺²
SevMos-C3 ⁴	0.375 ^{+2,3}	0.621 ^{+2,=3}	0.640 ^{+2,=3}
Delta-32 ⁵	0.382 ⁺³	0.629 ⁼³	0.648 ⁼³
Delta-32-Lex3	0.413^{+4,5}	0.638^{+4,5}	0.666^{+4,5}

Table 4: Ranking metrics on the ‘OneNewWord’ test data. The significance for DRMM’s NDCG.20 comparison is to a confidence of 95%.

Among the Delta models, Delta-32 showed better performance overall than SevMos-C1. However compared to SevMos-C3, its NDCG.20 score was lower, while the MAP and Prec.5 scores were higher. The Delta-32-Lex3 model exhibited the best metrics overall, bettering both Delta-32 and SevMos-C3. These gains were observed not just in the relevance-weighted NDCG.20 metrics that use our derived scaled relevance levels, but also in the MAP and Prec.5 precision-based metrics that use a binary notion of relevance.

The good performance of SevMos-C3 over SevMos-C1 demonstrates the benefits of using a convolutional stack. Combining these elements with the Delta matrix implementation of the Local Interaction architecture yields even better results, as depicted in the metrics for Delta-32-Lex3.

To evaluate performance on the *Under-Specified Query Problem*, we compared metrics on the ‘Neg20+’ subset of the test data (table 3). These are harder queries to rank for, since many non-relevant documents contain all the query words. Not surprisingly, the scores for all the models dropped. WMD was still the benchmark among untrained models, and DRMM the lowest performing deep learning model, although it did have better NDCG.20 score than BM25. The Delta-32-Lex3 model again exhibited the best overall performance.

To evaluate model robustness, we looked at performance on the ‘OneNewWord’ and ‘AllNewWords’ subsets of the test data (tables 4 and 5). The general trend among models here was the same as for the full test data, with DRMM performing no better than BM25, and SevMos-C3 better than SevMos-C1 and WMD. The Delta-32-Lex3 model showed the best overall performance, demonstrating that

	NDCG.20	MAP	Prec.5
rev DocID	0.195	0.508	0.389
BM25-Title ¹	0.309	0.581	0.586
WMD-Title ²	0.306 ⁼¹	0.590 ⁺¹	0.595 ⁼¹
DRMM	0.311 ^{=1,2}	0.578 ^{=1,-2}	0.570 ^{=1,-2}
SevMos-C1 ³	0.344 ⁺²	0.615 ⁺²	0.624 ⁺²
SevMos-C3 ⁴	0.355 ^{+2,=3}	0.614 ^{+2,=3}	0.632 ^{+2,=3}
Delta-32 ⁵	0.362 ^{+3,=4}	0.622 ^{+3,4}	0.638 ^{=3,4}
Delta-32-Lex3	0.400^{+4,5}	0.633^{+4,5}	0.661^{+4,5}

Table 5: Ranking metrics on the ‘AllNewWords’ test data

Figure 3: Comparing the impact of number of filters n_f (feature maps) in the Convolutional Stage of the Delta Model.

it was the most robust approach among the models tested. As a final note, some of the models exhibited better metrics on these subsets than the overall data because they tend to do better on shorter queries, and these test subsets had a higher concentration of shorter queries.

5.2 Impact of Different Features

Next we look at how different aspects of the Delta model affected its performance. The Convolutional stage extracts match-related features from word n -grams in the document. The parameter n_f controls the number of such features, and its effect on ranking is charted in figure 3. Both NDCG.20 and MAP improved as n_f increased till around 32 filters, and then performance leveled off and then dropped slightly. At larger number of filters, the model becomes more complex, but this increase in complexity does not continue to yield better performance. More complex models are more likely to overfit the training data, and perhaps learning rate decay techniques might help converge to a better solution, an area to be explored further. However since our goal was to construct a fast model for use in a search engine, we had a preference for smaller models, and $n_f = 32$ provided a good balance between speed and performance.

Table 6 compares four different versions of the Delta model with $n_f = 32$. The Delta Stage of the model computes, for each

	NDCG.20	MAP
Delta-32, no Difference vectors	0.323	0.574
Delta-32, no Delta features	0.333	0.584
Delta-32	0.365	0.601
Delta-32-Lex3, (top 3 Lex features)	0.394	0.609

Table 6: Comparing the impact of different features on the ranking metrics for the Delta model

document word, a difference vector against the closest query word, and three ‘Delta features’: the cosine similarity, euclidean distance, and normalized proximity. The first two rows of table 6 show the impact of removing the difference vectors and the Delta features from the Delta-32 model, on the ranking metrics for the test data. Both show a significant drop in performance compared to Delta-32. Finally, as reviewed above, adding the 3 lexical match features resulted in a significant improvement for the Delta-32-Lex3 model over the Delta-32 model.

In our greedy search for selecting the most useful lexical match features, BM25 on Document Abstract showed the most impact because it compensated for the truncation of the abstract when the document was limited to 50 words. The other two lexical match features are useful in accounting for the matches of out-of-vocabulary query terms (as discussed in [35]), and also for providing for query term significance through IDF-weighting.

5.3 Evaluation as a Reranker

Our goal for this research was to develop a good textual relevance model whose output could be used as a factor in a reranker (which would also use other factors like document meta-data) in a search engine like PubMed. We set a target of re-ranking the top 500 documents as produced by a fast naive ranker, and a run-time performance constraint on the relevance model of scoring the 500 documents in under 0.1 seconds on a GPU, yielding a throughput of at least 10 queries per second per GPU. While the complete design of a two-round ranker was outside the scope of this project, we compared the candidate relevance models on the (up to) top 500 results for the same test queries, as ranked by PubMed’s implementation of BM25, which also incorporates query expansion terms.

5.3.1 Ranking Metrics. Table 7 shows the ranking metrics for our candidate and baseline models on the top 500 results test data, with models sorted on the NDCG.20 metric. This data provided a particular challenge for relevance models, since on average less than 4% of the documents were relevant to the corresponding query. It also did not have the same selection bias present in the training data which was extracted from clicks on results sorted by PubMed’s relevance ranker, a more complex expression that includes BM25 as one of the factors. As a result, all the metrics were lower than for the previous test data. However the general trends among the models was the same, with Delta-32-Lex3 the best model overall by a large margin, followed by SevMos-C3. An area worth exploring further is whether adding some randomly sampled documents from deep in the search results (e.g. [17]) could help overcome some of this selection bias.

	NDCG.20	MAP
Original sort order	0.025	0.109
BM25-Title	0.032	0.077
SevMos-C1	0.083	0.104
DRMM	0.106	0.135
WMD-Title	0.122	0.153
Delta-32	0.141	0.154
SevMos-C3	0.160	0.151
Delta-32-Lex3	0.191	0.188

Table 7: Ranking metrics for re-ranking the top 500 results as provided by PubMed’s BM25, sorted on NDCG.20

5.3.2 Run-time Cost. As discussed above, the Independent Semantic Vector approach like that used in the SevMos models [35] is particularly attractive for use in search engines, because the document semantic vectors (e.g. pooled output from the convolution stage in SevMos) can be pre-computed and stored, the query vector needs to be computed just once per search, and the remaining computation (the similarity measure and classifier stage in SevMos) is fairly small and fast. This caching is not possible in the Local Interaction approaches like the DRMM and Delta models. In these two models, the computation is dominated by the cost to compare each pair of document and query words, so we look to reducing the size of the document by truncation to control the computation. In scientific literature, authors are strongly motivated to provide a highly informative and noise-free document Title and Abstract, making it particularly amenable to this approach.

We measured the time each model took to score 500 documents for a single query on a NVIDIA GeForce GTX TITAN X GPU, including the time to transfer the data to the GPU in a single batch, but not including the time to load the model and its parameters (including embeddings). While best ranking metrics for DRMM were obtained for $N = 200$ (with $\text{NDCG.20} = 0.300$), its run-time of 0.177 seconds, corresponding to a throughput of 5.6 queries per second per GPU, exceeded our criterion; DRMM for $N = 100$ came in at 0.079 seconds (throughput 12.7 qps per GPU), but a slightly lower overall NDCG.20 of 0.293 (figure 2b). The ‘Delta-32-Lex3’ model had a run-time of 0.049 secs, (20 qps per GPU). As a comparison the full ‘SevMos-C3’ model scored 500 documents in 0.040 seconds (25 qps per GPU).

5.4 Example Queries

We compare the rankings of the ‘Delta-32-Lex3’, ‘SevMos-C3’ and ‘WMD’ relevance models on some interesting queries. The NDCG at 20 at that query is quoted for each model, along with Titles and relevance levels for the top 3 scoring documents, and the relevance levels for the next 4 documents. The examples demonstrate the Delta model’s ability to detect relevance in documents without exact match of query terms (addressing the *term mismatch problem*), and where the context of the match is also important.

5.4.1 Query: countermovement jump. The word *countermovement* did not occur in training or validation queries. This is also an

example where relevance depends on the other words in the document text besides those matching the query. Number of documents in the test dataset: relevant = 17, non-relevant = 16. Top relevance levels: 21.5, 18.1, 11.2, 7.8, 4.4.

As ranked by **Delta-32-Lex3** (NDCG.20 = 0.98):

- i. (21.5) Determinants of countermovement jump performance: a kinetic and kinematic analysis.
- ii. (11.2) Which drop jump technique is most effective at enhancing countermovement jump ability, "countermovement" drop jump or "bounce" drop jump?
- iii. (4.4) The MARS for squat, countermovement, and standing long jump performance analyses: are measures reproducible?

Next four relevance levels = 4.4, 0, 18.1, 4.4.

As ranked by **SevMos-C3** (NDCG.20 = 0.32):

- i. (11.2) Which drop jump technique is most effective at enhancing countermovement jump ability, "countermovement" drop jump or "bounce" drop jump?
- ii. (0) A mechanics comparison between landing from a countermovement jump and landing from stepping off a box.
- iii. (4.4) Comparison of acute countermovement jump responses after functional isometric and dynamic half squats.

Next four relevance levels = 0, 4.4, 0, 21.5.

As ranked by **WMD** (NDCG.20 = 0.5):

- i. (11.2) Which drop jump technique is most effective at enhancing countermovement jump ability, "countermovement" drop jump or "bounce" drop jump?
- ii. (0) Reductions in Sprint Paddling Ability and Countermovement Jump Performance After Surfing Training.
- iii. (21.5) Determinants of countermovement jump performance: a kinetic and kinematic analysis.

Next four relevance levels = 4.4, 0, 0, 0.

5.4.2 Query: oesophageal cancer review. The word *oesophageal* did not occur in training or validation queries. The word *review* does not occur in the title of all relevant documents. The three models successfully located alternative spellings of the word. Number of documents in the test dataset: relevant = 22, non-relevant = 28. Top relevance levels: 11.7, and four at 7.1.

As ranked by **Delta-32-Lex3** (NDCG.20 = 0.94):

- i. (11.7) Esophageal cancer: Recent advances in screening, targeted therapy, and management.
- ii. (0) Outcomes in the management of esophageal cancer.
- iii. (0) Current advances in esophageal cancer proteomics.

Next four relevance levels = 5.6, 5.6, 0, 7.1.

As ranked by **SevMos-C3** (NDCG.20 = 0.05):

- i. (5.6) Esophageal Cancer Staging.
- ii. (2.5) Esophageal cancer: staging system and guidelines for staging and treatment.
- iii. (0) Outcomes in the management of esophageal cancer.

Next four relevance levels = 0, 0, 5.6, 0.

As ranked by **WMD** (NDCG.20 = 0.05):

- i. (5.6) Esophageal Cancer Staging.
- ii. (7.1) Endoscopic Management of Early Esophageal Cancer.
- iii. (0) Endoscopic treatment of early esophageal cancer.

Next four relevance levels = 0, 0, 5.6, 0.

5.4.3 Query: chronic headache and depression review. All three models were able to leverage word vectors to relate headache to migraine. Delta-32-Lex3 placed the most relevant document ("Psychological Risk Factors in Headache", rel. level = 10) at rank 5. It did not feature in the top 10 of any of the other neural and lexical models tested. This example demonstrates the need for deeper semantic modeling, where the Delta model has some limited success. Number of documents in the test dataset: relevant = 23, non-relevant = 18. Top relevance levels: a 10, and four at 5.5.

As ranked by **Delta-32-Lex3** (NDCG.20 = 0.4):

- i. (5.5) Comprehensive management of headache and depression.
- ii. (0) Medication overuse headache.
- iii. (0) Clinical features and mechanisms of chronic migraine and medication-overuse headache.

Next four relevance levels = 0, 10, 5.5, 5.5.

As ranked by **SevMos-C3** (NDCG.20 = 0.3):

- i. (5.5) Chronic headaches and the neurobiology of somatization.
- ii. (2.7) Pathophysiology of migraine – from molecular to personalized medicine.
- iii. (0) Medication overuse headache.

Next four relevance levels = 5.5, 2.5, 5.5, 0.

As ranked by **WMD** (NDCG.20 = 0.35):

- i. (5.5) Comprehensive management of headache and depression.
- ii. (5.5) Migraine and depression: biological aspects.
- iii. (5.5) Migraine and depression.

Next four relevance levels = 5.5, 0, 0, 0.

6 CONCLUSION

While deep learning models for text understanding have made dramatic gains in recent years, they have tended to be large and slow. The challenge in information retrieval is still one of combining predictive power with low run-time overhead. This is also true when the corpus is scientific literature.

We described the Delta Relevance model, a new deep learning model for text relevance, targeted for information retrieval in biomedical science literature. The main innovation in the model is to base the modeling function on differences and distances between word embeddings, as captured in the Delta features, rather than directly using the embeddings themselves like most NLP approaches. Other researchers have shown the benefit of training word embeddings as part of the model. That was not feasible here since the training data had only 20k queries compared to a vocabulary size of 200k. Using Delta features as the input to the model helped adapt the pre-trained word embeddings to our task.

To achieve our goal of fast run-time, we used a convolutional network rather than the recurrent approaches popular in most neural NLP models. Using a stack of narrow convolutional layers instead of a single wide convolution gave the model more power.

We showed that the Delta model outperformed comparable recent approaches in ranking metrics when trained and evaluated on data derived from the PubMed search engine click logs. We demonstrated that the model was robust, despite being trained on a relatively small amount of data, and the model was fast enough for use in an on-line search engine.

The Delta model might be especially suited to scientific literature in taking advantage of the high quality of the document Title and first few sentences of the Abstract. We believe that the previously observed good performance of the DRMM approach was on documents that were quite noisy (they contain a lot of meta data in the document text). An area worth exploring further, for its potential in improving both prediction performance and run-time costs, is pre-processing a document to extract significant portions for evaluating relevance, thus reducing the size of the input at run-time. Another area to explore is the benefit from retaining sentence and grammatical structure in document text.

ACKNOWLEDGMENTS

This research was supported by the Intramural Research Program of the NIH, National Library of Medicine.

REFERENCES

- [1] Olivier Bodenreider. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. 32 (Database issue) (2004), D267–D270.
- [2] Chris Buckley and Ellen M. Voorhees. 2000. Evaluating Evaluation Measure Stability. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '00)*. ACM, New York, NY, USA, 33–40.
- [3] Billy Chiu, Gamal Crichton, Anna Korhonen, and Sampo Pyysalo. 2016. How to Train good Word Embeddings for Biomedical NLP. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing (BioNLP '16)*. ACL, 166–174.
- [4] Fernando Diaz, Bhaskar Mitra, and Nick Craswell. 2016. Query Expansion with Locally-Trained Word Embeddings. In *Proceedings of ACL 2016*. ACL, Berlin, Germany, 367–377.
- [5] Rezarta Islamaj Dogan, G. Craig Murray, Aurélie Névél, and Zhiyong Lu. 2009. Understanding PubMed® user search behavior through log analysis. *Database* 2009 (2009), bap018.
- [6] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12 (Jul 2011), 2121–2159.
- [7] Frederico Durao, Karunakar Bayyapu, Guandong Xu, Peter Dolog, and Ricardo Lage. 2013. *Medical Information Retrieval Enhanced with User's Query Expanded with Tag-Neighbors*. Springer New York, New York, NY, 17–40. https://doi.org/10.1007/978-1-4614-8495-0_2
- [8] G. W. Furnas, T. K. Landauer, L. M. Gomez, and S. T. Dumais. 1987. The Vocabulary Problem in Human system Communication. *CACM* 30, 11 (nov 1987), 964–971.
- [9] Jianfeng Gao, Patrick Pantel, Michael Gamon, Xiaodong He, and Li Deng. 2014. Modeling Interestingness with Deep Neural Networks. In *Proceedings of EMNLP 2014*. ACL, Doha, Qatar, 2–13.
- [10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press, Cambridge, MA. <http://www.deeplearningbook.org>.
- [11] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. A Deep Relevance Matching Model for Ad-hoc Retrieval. In *Proceedings of CIKM 2016*. ACM, New York, NY, USA, 55–64.
- [12] W. Hersch, S. Price, and L. Donohoe. 2000. Assessing Thesaurus-Based Query Expansion Using the UMLS Metathesaurus. In *Proceedings of the AMIA Symposium*. 344–348.
- [13] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional Neural Network Architectures for Matching Natural Language Sentences. In *Advances in NIPS 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Eds.). 2042–2050.
- [14] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of CIKM 2013*. ACM, New York, NY, USA, 2333–2338.
- [15] Kai Hui, Andrew Yates, Klaus Berberich, and Gerard de Melo. 2017. PACRR: A Position-Aware Neural IR Model for Relevance Matching. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1060–1069. <http://aclweb.org/anthology/D17-1111>
- [16] Kalervo Järvelin and Jaana Kekäläinen. 2000. IR Evaluation Methods for Retrieving Highly Relevant Documents. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '00)*. ACM, New York, NY, USA, 41–48.
- [17] Thorsten Joachims. 2002. Optimizing Search Engines Using Clickthrough Data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '02)*. ACM, New York, NY, USA, 133–142.
- [18] Sun Kim, Zhiyong Lu, and W. John Wilbur. 2015. Identifying named entities from PubMedÂ for enriching semantic categories. *BMC Bioinformatics* 16, 1 (21 Feb 2015), 57. <https://doi.org/10.1186/s12859-015-0487-2>
- [19] Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From Word Embeddings To Document Distances. In *Proceedings of The 32nd International Conference on Machine Learning (ICML 2015)*. JMLR, 957â–966.
- [20] John Lafferty and Chengxiang Zhai. 2001. Document Language Models, Query Models, and Risk Minimization for Information Retrieval. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '01)*. ACM, New York, NY, USA, 111–119. <https://doi.org/10.1145/383952.383970>
- [21] Yann LeCun. 1989. Generalization and Network Design Strategies. In *Connectionism in Perspective*, R. Pfeifer, Z. Schreier, F. Fogelman, and L. Steels (Eds.). Elsevier, Zurich, Switzerland.
- [22] Zhiyong Lu, Won Kim, and W. John Wilbur. 2009. Evaluation of query expansion using MeSH in PubMed. *Information retrieval* 12, 1 (May 2009), 69–80.
- [23] Zhengdong Lu and Hang Li. 2013. A Deep Architecture for Matching Short Texts. In *Advances in NIPS 26*, C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). 1367–1375.
- [24] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. 2013. Rectifier Nonlinearities Improve Neural Network Acoustic Models. In *Proceedings of the ICML Workshop on Deep Learning for Audio, Speech, and Language Processing (WDLASL 2013)*.
- [25] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *Proceedings of the Workshop at ICLR 2013*.
- [26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in NIPS 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). 3111–3119.
- [27] David R. H. Miller, Tim Leek, and Richard M. Schwartz. 1999. A Hidden Markov Model Information Retrieval System. In *Proceedings of SIGIR 1999*. ACM, New York, NY, USA, 214–221.
- [28] Bhaskar Mitra and Nick Craswell. 2017. Neural Models for Information Retrieval. *CoRR* abs/1705.01509 (2017). <https://arxiv.org/abs/1705.01509>
- [29] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to Match Using Local and Distributed Representations of Text for Web Search. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1291–1299. <https://doi.org/10.1145/3038912.3052579>
- [30] Sunil Mohan, Nicolas Fiorini, Sun Kim, and Zhiyong Lu. 2017. Deep Learning for Biomedical Information Retrieval: Learning Textual Relevance from Click Logs. In *BioNLP 2017*. Association for Computational Linguistics, Vancouver, Canada, 222–231. <http://www.aclweb.org/anthology/W17-2328>
- [31] Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. 2016. Improving Document Ranking with Dual Word Embeddings. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 83–84. <https://doi.org/10.1145/2872518.2889361>
- [32] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text Matching as Image Recognition. (2016). <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11895>
- [33] Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1994. Okapi at TREC-3. In *Proceedings of TREC 1994*. NIST, Dept. of Commerce, 109–126.
- [34] Y. Rubner, C. Tomasi, and L. J. Guibas. 1998. A metric for distributions with applications to image databases. In *Proceedings of the Sixth International Conference on Computer Vision*. IEEE.
- [35] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to Rank Short Text Pairs with Convolutional Deep Neural Networks. In *Proceedings of SIGIR 2015*. ACM, New York, NY, USA, 373–382.
- [36] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. A Latent Semantic Model with Convolutional-Pooling Structure for Information Retrieval. In *Proceedings of CIKM 2014*. ACM, New York, NY, USA, 101–110.
- [37] Gregor Urban, Krzysztof J. Geras, Samira Ebrahimi Kahou, Ozlem Aslan, Shengjie Wang, Abdelrahman Mohamed, Matthai Philipose, Matt Richardson, and Rich Caruana. 2017. Do Deep Convolutional Nets Really Need to be Deep and Convolutional?. In *Proceedings of the Fifth International Conference on Learning Representations (ICLR 2017)*. Toulon, France.
- [38] Chenyan Xiong, Zhuyun Dai, Jamie Callan, Zhiyuan Liu, and Russell Power. 2017. End-to-End Neural Ad-hoc Ranking with Kernel Pooling. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '17)*. ACM, New York, NY, USA, 55–64. <https://doi.org/10.1145/3077136.3080809>
- [39] Chengxiang Zhai and John Lafferty. 2004. A Study of Smoothing Methods for Language Models Applied to Information Retrieval. *ACM Trans. Inf. Syst.* 22, 2 (April 2004), 179–214.
- [40] Y. Zhang, M. M. Rahman, A. Braylan, B. Dang, H. Chang, H. Kim, Q. McNamara, A. Angert, E. Banner, V. Khetan, T. McDonnell, A. T. Nguyen, D. Xu, B. C. Wallace, and M. Lease. 2016. Neural Information Retrieval: A Literature Review. *CoRR* abs/1611.06792 (2016). <http://arxiv.org/abs/1611.06792>