A Feedback-Based Approach to Utilizing Embeddings for Clinical Decision Support

Chenhao Yang¹ · Ben He¹ · Canjia Li¹ · Jungang Xu¹

Received: 24 August 2017 / Revised: 18 October 2017 / Accepted: 31 October 2017 / Published online: 10 November 2017
© The Author(s) 2017. This article is an open access publication

Abstract Clinical Decision Support (CDS) is widely seen as an information retrieval (IR) application in the medical domain. The goal of CDS is to help physicians find useful information from a collection of medical articles with respect to the given patient records, in order to take the best care of their patients. Most of the existing CDS methods do not sufficiently consider the semantic relation between texts, hence the potential in improving the performance in biomedical articles retrieval. This paper proposes a novel feedback-based approach which considers the semantic association between a retrieved biomedical article and a pseudo feedback set. Evaluation results show that our method outperforms the strong baselines and is able to improve over the best runs in the TREC CDS tasks.

Keywords Clinical Decision Support · Semantic association · Relevance feedback

1 Introduction

The goal of Clinical Decision Support (CDS) is to efficiently and effectively link relevant biomedical articles to meet physicians' needs for taking better care of their patients. In CDS applications, the patient records are considered as queries and the biomedical articles are retrieved in response to the queries. A major difference between CDS and traditional IR tasks is that the documents, mostly scientific articles, are very long and contain comprehensive

information about a specific topic such as a treatment for a disease, or a patient case. As a result, the CDS queries, although longer than those in other IR tasks, may not cover the various aspects of the user information need, and simple document-query matching does not lead to optimal effectiveness in the CDS task.

Most of the existing CDS methods retrieve biomedical articles using the frequency-based statistical models [1, 2, 6, 9]. Those methods extract concepts from queries and biomedical articles, and further utilize concepts to apply query expansion or document ranking. Then, the relevance score of a given article is assigned based on the frequencies of query terms or concepts. Despite the fact that the frequency-based CDS methods have been shown to be effective and efficient in the CDS task [25], they ignore the semantic associations between texts. We argue that the retrieval effectiveness of the CDS systems can be further improved by integrating the semantic information. For instance, suppose two short medical-related texts as follows:

- The child has symptoms of strawberry red tongue and swollen red hands.
- This kid is suffering from Kawasaki disease.

Though the two short sentences have no terms in common, they convey the same meaning and are considered to be related to each other. However, the two sentences above are considered completely unrelated by the existing frequency-based CDS methods. In this paper, we aim to further enhance the retrieval performance of the CDS systems by taking the semantic association between texts into consideration. Benefiting from recent advances in natural language processing (NLP), words and documents can be represented with semantically distributed real-valued vectors, namely the *embeddings*, which are generated by

✉ Ben He
benhe@ucas.ac.cn

¹ University of the Chinese Academy of Sciences, Beijing, China

neural network models [3, 17, 21, 22]. The embeddings have been shown to be effective and efficient in many NLP tasks due to the ability in preserving semantic relationships in vector operations such as summation and subtraction [21]. In this study, we utilize the Word2Vec technique proposed by Mikolov et al. [17, 21] to generate embeddings of words and biomedical articles, which is widely considered as an effective embedding method in NLP applications [8, 20, 30]. As a state-of-the-art topic model, latent Dirichlet allocation (LDA) [5] is also used for comparison with Word2Vec in generating distributed representations of biomedical articles in this study.

There have been efforts in utilizing the embeddings to improve IR effectiveness. For example, Vulić and Moens estimate a semantic relevance score by the cosine similarity between the embeddings of the query–document pair to improve the performance of monolingual and cross-lingual retrieval [31]. Similar idea is presented in [32], where the semantic similarity between the embeddings of the patient record and biomedical article is utilized to improve the CDS system. We argue that query is a weak indicator of relevance in that query is usually much shorter than the relevant documents, such that the use of semantic associations of the query–document pairs may only lead to limited improvement in retrieval performance. To this end, this paper proposes a feedback-based CDS method which integrates semantic associations between texts to further enhance retrieval effectiveness. To the best of our knowledge, this paper is the first to estimate the relevance score for IR tasks based on document-to-document (D2D) embedding similarity. Experimental results show that our proposed CDS method can have significant improvements over strong baselines. In particular, a simple linear combination of the classical BM25 weighting function with the semantic relevance score generated by our method leads to effective retrieval results that are better than the best TREC CDS runs.

A conference version of this paper was published in [33]. Extensions to the conference version include:

- The experiments conducted on the recent TREC 2016 CDS task dataset. The results obtained on this new dataset are consistent with those on the TREC 2014 and 2015 CDS datasets.
- The proposed approach is further evaluated on five standard IR test collections. Results show that our approach is able to outperform strong baselines for IR tasks other than CDS.

The remainder of this paper is organized as follows. Section 2 briefly introduces the related work. Section 3 describes the proposed feedback-based approach in details. For the evaluation of the proposed approach on the CDS datasets, the experimental settings and results are presented

in Sects. 4 and 5, respectively. The proposed approach is further evaluated on other standard TREC IR collections in Sect. 6. Finally, Sect. 7 concludes this work and suggests possible future research directions.

2 Related Work

2.1 BM25 and PRF

As our CDS method is to integrate the semantic relevance score into the classical BM25 model with applying pseudo-relevance feedback (PRF), we introduce BM25 model and PRF in this section. The ranking function of BM25 given a query Q and a document d is as follows [26]:

$$\text{score}(d, Q) = \sum_{t \in Q} w_t \frac{(k_1 + 1)tf}{K + tf} \frac{(k_3 + 1)qtf}{k_3 + qtf} \quad (1)$$

where t is one of the query terms, and qtf is the frequency of t in query Q . tf is the term frequency of query term t in document d . K is given by $k_1((1 - b) + b \cdot \frac{l}{\text{avg}_l})$, in which l and avg_l denote the length of document d and the average length of documents in the whole collection, respectively. k_1 , k_3 and b are free parameters whose default setting is $k_1 = 1.2$, $k_3 = 1000$ and $b = 0.75$, respectively [26]. w_t is the weight of query term t , which is given by:

$$w_t = \log_2 \frac{N - df_t + 0.5}{df_t + 0.5} \quad (2)$$

where N is the number of documents in the collection, and df_t is the document frequency of query term t , which denotes the number of documents that t occurs.

Pseudo-relevance feedback (PRF) is a popular method for improving IR effectiveness by using the top- k documents as pseudo-relevance set [18]. One of the best-performing PRF methods on top of BM25 is an adoption of Rocchio's algorithm presented in [16], which is able to provide state-of-the-art retrieval effectiveness on standard TREC test collections [16]. BM25 with PRF is denoted as $BM25_{PRF}$ in this paper.

2.2 State-of-the-Art CDS Methods

Due to the specificity of medical healthcare field, most of the existing CDS methods retrieve biomedical articles based on concepts, including unigrams, bigrams and multi-word concepts. These concepts are extracted from different resources, such as queries, biomedical articles, external medical databases, etc. These content-based CDS methods usually utilize concepts to apply query expansion or document ranking based on the frequencies of the concepts. Palotti and Hanbury proposed a concept-based query

expansion method, increasing the weights of relevant concepts and expanding the original query with concepts extracted by MetaMap [23]. MetaMap is a highly configurable tool for recognizing the Unified Medical Language System (UMLS) concepts in text, which is usually utilized in the existing CDS methods. Song et al. proposed a customized learning-to-rank algorithm and a query term position-based re-ranking model to improve the retrieval performance [28]. As biomedical articles are usually full-text scientific articles which are much longer than Web documents, Cummins et al. applied the recently proposed SPUD language model [10] to CDS for retrieving long documents in a balanced way [9]. Abacha and Khelifi investigated several query reformulation methods utilizing Mesh and DBpedia. In addition, they applied rank fusion to combine different ranked document lists into a single list to improve the retrieval performance [1].

2.3 The Best-Performing Methods in the TREC CDS Tasks

Choi and Choi proposed a three-step biomedical article retrieval method, which obtains the best run in the TREC 2014 CDS task [6]. Firstly, the method utilizes external knowledge resource to apply query expansion, and uses the query likelihood (QL) language model [24] to rank articles. Secondly, a text classification based method is used for the topic-specific ranking. Note that the topics used in the TREC CDS task are classified into three categories, i.e., *diagnosis*, *test* and *treatment*. Finally, the method combines the relevance ranking score and the topic-specific ranking score with Borda-fuse method [6].

The CDS methods proposed by Balaneshin-kordan et al. [2] obtained both the best automatic and manual runs in the TREC 2015 CDS task. Their method extracts unigrams, bigrams and multi-word UMLS concepts from queries, the pseudo-relevance feedback documents or external knowledge resources, and then uses the Markov Random Field (MRF) model [19] for document ranking. The relevance score of a document d given a query Q is computed as follows [2]:

$$\begin{aligned} score(d, Q) &= \sum_{c \in \mathbb{C}} \mathbb{1}_c score(c, d) \\ &= \sum_{c \in \mathbb{C}} \mathbb{1}_c \sum_{T \in \mathbb{T}} \lambda_T f_T(c, d) \end{aligned} \quad (3)$$

where $score(c, d)$ is the contribution of concept c to the relevance score of document d . $\mathbb{1}_c$ is an indicator function which determines whether the concept c is considered in the relevance weighting. $\mathbb{C}$ is the set of concepts. $\mathbb{T}$ is the set of all concept types, to which concept c belongs. Note that a concept can belong to multiple concept types at the

same time. λ_T is the importance weight of concept type T , and $f_T(c, d)$ is a real-valued feature function.

In the TREC 2016 CDS task, Harsha Gurulingappa et al. [15] proposed a semi-supervised method that takes the advantages of pseudo-relevance feedback, semantic query expansion and document similarity measures based on unsupervised word embeddings. Firstly, terms expanded by the UMLS concepts and document titles in the top- k pseudo-relevance feedback set are extracted and added with a weight of 0.1 to the initial query. Secondly, by using the unsupervised word embedding method, centroids for articles are computed based on the abstract, the title or the journal title. Finally, ranking scores obtained from PRF, UMLS expansion, and word embedding document distances are used as features to the supervised learning to rank model.

The existing CDS methods retrieve biomedical articles based on the frequencies of concepts. As discussed in Sect. 1, the lack of semantic associations between texts may lead to limited retrieval performance. A recent work [32] integrates semantic similarity between the embeddings of the patient record and biomedical article to improve the CDS system, which is given by:

$$Sim(d, Q) = 0.5 \cdot \frac{\vec{d} \cdot \vec{Q}}{\|\vec{d}\| \times \|\vec{Q}\|} + 0.5 \quad (4)$$

where $\vec{d}$ and $\vec{Q}$ are the embeddings of biomedical article d and patient record Q , respectively. $Sim(d, Q)$ is the semantic similarity which is integrated into BM25 model [26] by a linear interpolation. As the patient records are usually much shorter than the full-text biomedical articles, they do not necessarily contain sufficient amount of semantic evidence of relevance. Therefore, the approach in [32] leads to limited improvement on the CDS task. To deal with this problem, in the next section, we propose a feedback-based approach that considers the semantic similarity between a retrieved article and a set of feedback articles, which is a better indicator of relevance than patient record.

3 Feedback-Based Semantic Relevance

The methods for generating the embeddings of biomedical articles are introduced in Sect. 3.1. The generated embeddings are utilized for enhancing the retrieval performance of CDS in Sect. 3.2.

3.1 Generating Embeddings of Biomedical Articles

The Word2Vec technique proposed by Mikolov et al. [17, 21] is a state-of-the-art neural embedding framework,

which has been shown to be effective and efficient in many NLP tasks. In this study, Word2Vec is also utilized to generate embeddings of words and biomedical articles. A unique advantage of Word2Vec is that the semantic relationships can be preserved in vector operations, such as addition and subtraction [21]. Therefore, the embeddings of biomedical articles can be generated through vector operations of word embeddings such that they are applicable to the CDS task. Considering the fact that informative words are usually infrequent in biomedical articles, we utilize the Skip-gram architecture of Word2Vec, which shows better performance for infrequent words than the CBOW architecture of Word2Vec in generating embeddings [21]. Besides, the negative sampling algorithm is used to train embeddings [21].

The Skip-gram architecture is composed of three layers, i.e., an input layer, a projection layer and an output layer. The basic idea of Skip-gram is to predict the context of a given word w . Given a word w and the corresponding context $c(w)$, the conditional probability $p(c(w)|w)$ is modeled by Softmax regression, which is given as follows:

$$p(c(w)|w; \theta) = \frac{e^{v_w \cdot v_{c(w)}}}{\sum_{c'(w) \in C} e^{v_w \cdot v_{c'(w)}}} \quad (5)$$

where v_w and $v_{c(w)}$ are the embeddings of word w and the corresponding context $c(w)$, respectively. The goal of Skip-gram model is to maximize the likelihood function of Equation 5 as follows [13]:

$$\begin{aligned} \arg \max_{\theta} \prod_{(w, c(w)) \in D} \log p(c(w)|w) \\ = \sum_{(w, c(w)) \in D} \left(\log e^{v_w \cdot v_{c(w)}} - \log \sum_{c'} e^{v_w \cdot v_{c'(w)}} \right) \end{aligned} \quad (6)$$

where w and $c(w)$ denote a word and the corresponding context, respectively. $(w, c(w))$ is a training sample, and D is the set of all training samples. θ is the parameter set that trained by stochastic gradient ascent.

A major challenge of the application of the word embeddings to CDS is how to generate effective embeddings for biomedical articles. In this paper, we adopt two ways of generating embeddings for biomedical articles, namely *Term Summation* and *Paragraph Embeddings*, abbreviated as *Sum* and *Para*, respectively. As the semantic relationships are preserved in the embedding operations, one way of generating embeddings of biomedical articles is to sum up the word embeddings of the top- k most informative words in a given article, i.e., *Term Summation*, which is given by:

$$\vec{d} = \sum_{w \in W_k^d} tf-idf(w) \cdot \vec{w} \quad (7)$$

where $\vec{w}$ and $\vec{d}$ are the embeddings of word w and biomedical article d , respectively. W_k^d is the set of the top- k terms with the highest $tf-idf$ weights in d . $tf-idf(w)$ is used to measure the amount of information carried by word w , which is given by:

$$tf-idf(w) = tf \cdot \log_2 \frac{N - df_w + 0.5}{df_w + 0.5} \quad (8)$$

where tf is the term frequency of w in d . N is the total number of biomedical articles in the whole collection, and df_w is the document frequency of word w .

In addition to *Term Summation*, we adopt the *Paragraph Embeddings* technique [17] to generate embeddings of biomedical articles. *Paragraph Embeddings* is an improved version of Word2Vec, in which each document is marked with a special word called *Paragraph id*. The *Paragraph id* participates in the training of each word as part of each context, acting as a memory that remembers what is missing from the current context. The training procedure of *Paragraph Embeddings* is the same as *Word2Vec*. Finally, the embedding of the special word *Paragraph id* is used to represent the corresponding biomedical article. We denote embeddings of biomedical articles generated by *Term Summation* and *Paragraph Embeddings* as $\vec{d}_{Sum}$ and $\vec{d}_{Para}$, respectively.

3.2 Using Embeddings for CDS

In this section, we introduce our proposed feedback-based CDS method, which considers the semantic similarity between a biomedical article to be scored and a pseudo feedback set. As Mikolov et al. demonstrated that words can have multiple degrees of similarity [21], integrating semantic associations by directly measuring the similarity between the embeddings of patient records and biomedical articles may only lead to limited improvement in retrieval performance (as used in [32]). Instead, we estimate the semantic relevance of a biomedical article by measuring the semantic similarity between the article and a pseudo-relevance feedback set. Once we obtain the preliminary retrieval results returned by BM25, the semantic relevance score of biomedical articles can be utilized to improve the retrieval performance, which is given as follows:

$$score(d, Q) = \lambda \cdot BM25(d, Q) + (1 - \lambda) \cdot SEM(d, D_{PRF}^k(Q)) \quad (9)$$

where $BM25(d, Q)$ is the ranking score of document d given by a baseline retrieval model, e.g., the classical BM25 model with PRF. $D_{PRF}^k(Q)$ is the pseudo-relevance feedback set of biomedical articles, which is composed of the top ranked k articles returned by the baseline model. It is usually assumed by the PRF technique that most of the

documents in $D_{PRF}^k(Q)$ are relevant to query Q , thus $D_{PRF}^k(Q)$ can be considered as a better indicator of relevance than patient records. $SEM(d, D_{PRF}^k(Q))$ measures the semantic similarity between document d and the pseudo-relevance feedback set $D_{PRF}^k(Q)$, which is given as follows:

$$SEM(d, D_{PRF}^k(Q)) = \sum_{d' \in D_{PRF}^k(Q)} w_{d'} \cdot Sim(d', d) \quad (10)$$

where d' is one of the biomedical articles in $D_{PRF}^k(Q)$. $w_{d'}$ is the importance weight of d' , which is given as follows:

$$w_{d'} = BM25(d', Q) + \max_{d'' \in D_{PRF}^k(Q)} BM25(d'', Q) \quad (11)$$

$Sim(d', d)$ denotes the semantic similarity between d' and d , which is given by Equation (4). In Equation (11), the maximum relevance score is added to normalize the gap between the relevance scores of different articles. Note that both $BM25(d, Q)$ and $SEM(d, D_{PRF}^k(Q))$ in Equation (9) are normalized by Min–Max normalization, such that the two scoring features are on the same scale.

4 Experimental Settings

In this section, we introduce the datasets used in the experiments and the experimental design.

4.1 Datasets

All our experiments are conducted on the standard datasets used in the TREC CDS tasks of 2014, 2015 and 2016. The target document collection used is an open access subset¹ of PubMed Central² (PMC). In 2014 and 2015, the same 733,138 articles were extracted, and in 2016, a larger and newer set of 1.25 million articles were extracted. We extract the *title*, *abstract*, *keywords* and *body* fields from each article as the source of the index. We use the open source Terrier toolkit version 4.1 [20] to index the collection with the recommended settings of the toolkit. Standard English stopwords are removed and the collection is stemmed using Porter's English stemmer. Using Porter's stemmer, inflected or derived words are reduced to their word stem, base or root forms.

There are 30 topics in each year, and each topic is a medical record narrative that serves as an idealized representation of actual patient record. These topics are classified into three categories, i.e., *diagnosis*, *test* and *treatment*, with 10 topics in each category. According to [27], there is little difference observed in retrieval performance when the

three topic types are taken into account. Thus the topic types are not considered in our study. For the previous two years, there are two versions of the medical record narratives, i.e., *Summary* and *Description* fields. The *Description* field is much longer than the *Summary* field, and has more detailed information about a patient. However, the *Description* field may contain more irrelevant information than the *Summary* field. In 2016, *Note* field was added into the topics, which contains patients chief complaint, relevant medical history and lots of other necessary information. Table 1 presents an example of the *Summary*, *Description* and *Note* fields. In the experiments, the *Summary*, *Description* and *Note* fields are separately used as queries.

As described in Sect. 3.1, the Skip-gram model of Word2Vec³ toolkit is utilized to generate embeddings of words and biomedical articles, which are trained using the negative sampling algorithm [21]. Note that the *title*, *abstract*, *keywords* and *body* fields of each biomedical article are extracted as the training set of Word2Vec, and the stopword removal and stemming are applied. As recommended in [21], the window size is set to 10 for Skip-gram model. As documents in the target collection are full-text long biomedical articles, the number of dimensions of the embeddings are set to 300, a value that is larger than the recommended 100 in [11].

4.2 Experimental Design

In our study, we evaluate our CDS method against two baselines. As described in Sect. 3.2, we use the BM25 model [26] with applying PRF as one of the baselines. In addition, we use the CDS method proposed in [32] as another baseline.

The parameters k_1 and k_3 of BM25 (See Equation (1)) are set to default values and b is set to the optimal value on training data by grid search algorithm [4]. As described in Sect. 3.1, we adopt two methods for generating embeddings of biomedical articles, which are denoted as $\overrightarrow{d_{Sum}}$ and $\overrightarrow{d_{Para}}$, respectively. For convenience, we denote our proposed CDS method applying *Term Summation* and *Paragraph Embeddings* as $BM25 + SEM_{d_{Sum}-D_{PRF}^k}$ and $BM25 + SEM_{d_{Para}-D_{PRF}^k}$, respectively. Besides, the previously proposed CDS method [32] is denoted as $BM25 + Sim_{d_{Para}-Q}$, which only uses *Paragraph Embeddings* for generating embeddings of biomedical articles. Note that our method has the following tunable parameters, i.e., hyper-parameter λ (see Equation (9)), top $|T|$ terms to generate embeddings of biomedical articles when applying

¹ <http://www.ncbi.nlm.nih.gov/pmc/tools/openftlist/>.

² <http://www.ncbi.nlm.nih.gov/pmc>.

³ The learned embeddings of words and biomedical articles can be downloaded from <http://gucasir.org/CDS.tgz>.

Table 1 Example of *Summary*, *Description* and *Note* fields**Topic type** - diagnosis**Summary:** A 78-year-old male presents with frequent stools and melena.**Description:** 78 M transferred to nursing home for rehab after CABG. Reportedly readmitted with a small NQWMI. Yesterday, he was noted to have a melanotic stool and then today he had approximately 9 loose BM some melena and some frank blood just prior to transfer, unclear quantity**Note:** 78 M w/pmh of CABG in early [*Month(only)*3 **] at [*Hospital*64406 **] (transferred to nursing home for rehab on [*12 - 8 ***] after several falls out of bed.) He was then readmitted to [*Hospital*61749 **] on [*3120 - 12 - 11 ***] after developing acute pulmonary edema/CHF/unresponsiveness?. There was a question whether he had a small MI; he reportedly had a small NQWMI. He improved with diuresis and was not intubated. Yesterday, he was noted to have a melanotic stool earlier this evening and then approximately 9 loose BM w/ some melena and some frank blood just prior to transfer, unclear quantity**Table 2** Evaluation results on the TREC 2014 CDS task

Model	infNDCG	infAP	MAP	R-Prec
Summary field				
BM25	0.2524	0.0805	0.1537	0.2004
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2698 + 6.89%*	0.0935 16.15%*	0.1628 + 5.92%*	0.2067 + 3.14%
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.2748 + 8.87%*	0.0953 + 18.39%*	0.1645 + 7.03%*	0.2083 + 3.94%
Description field				
BM25	0.2460	0.0700	0.1440	0.2065
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2751 + 11.83%*	0.0918 + 31.14%*	0.1623 + 12.71%*	0.2196 + 6.34%*
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.2830 + 15.04%*	0.0911 + 30.14%*	0.1661 + 15.35%*	0.2206 + 6.83%*

The difference in percentage is measured against the baseline retrieval model *BM25*. A statistically significant difference is marked with a *. The best result of each evaluation metric is in bold

Table 3 Evaluation results on the TREC 2015 CDS task

Model	infNDCG	infAP	MAP	R-Prec
Summary field				
BM25	0.2695	0.0736	0.1650	0.2198
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2980 + 10.58%*	0.0831 12.91%*	0.1758 + 6.55%*	0.2345 + 6.69%*
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.2986 + 10.80%*	0.0842 + 14.40%*	0.1791 + 8.55%*	0.2408 + 9.55%*
Description field				
BM25	0.2724	0.0733	0.1641	0.2184
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2877 + 5.62%*	0.0837 14.19%*	0.1762 + 7.37%*	0.2325 + 6.46%*
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.3016 + 10.72%*	0.0873 + 19.10%*	0.1806 + 10.05%*	0.2370 + 8.52%*

The difference in percentage is measured against the baseline retrieval model *BM25*. A statistically significant difference is marked with a *. The best result of each evaluation metric is in bold

Term Summation (# *Terms*) and top k articles in $D_{PRF}^k(Q)$ (# *PRF Documents*). All the parameters are tuned on training data by grid search algorithm [4].

The evaluation results are obtained by a twofold cross-validation, where the topics are split into two equal-size subsets by parity in odd or even topic numbers. In each fold, we use one subset of the topics for training, and the

Table 4 Evaluation results on the TREC 2016 CDS task

Model	infNDCG	infAP	MAP	R-Prec
Summary field				
$BM25_{PRF}$	0.2081	0.0274	0.0806	0.1491
$BM25_{PRF} + SEM_{d_{pr}-D_{PRF}^k}$	0.2260	0.0318	0.0817	0.1501
	+ 8.60%*	+ 16.06%*	+ 1.36%	+ 0.67%
$BM25_{PRF} + SEM_{d_{add}-D_{PRF}^k}$	0.2493	0.0345	0.0837	0.1486
	+ 19.80%*	+ 25.91%*	+ 3.85%*	-0.33%
Description field				
$BM25_{PRF}$	0.1547	0.0153	0.0523	0.1121
$BM25_{PRF} + SEM_{d_{pr}-D_{PRF}^k}$	0.1724	0.0222	0.0594	0.1148
	+ 11.44%*	+ 45.10%*	+ 13.58%*	+ 2.40%
$BM25_{PRF} + SEM_{d_{add}-D_{PRF}^k}$	0.1786	0.0225	0.0583	0.1195
	+ 15.45%*	+ 47.06%*	+ 11.47%*	+ 6.60%*
Note Field				
$BM25_{PRF}$	0.1698	0.0206	0.0669	0.1197
$BM25_{PRF} + SEM_{d_{pr}-D_{PRF}^k}$	0.1849	0.0242	0.0665	0.1154
	+ 8.90%*	+ 17.90%*	-0.60%	-3.60%
$BM25_{PRF} + SEM_{d_{add}-D_{PRF}^k}$	0.1957	0.0255	0.0709	0.1243
	+ 15.25%*	+ 23.79%*	+ 5.98%*	+ 3.84%*

The difference in percentage is measured against the baseline retrieval model $BM25$. A statistically significant difference is marked with a *. The best result of each evaluation metric is in bold

Table 5 Comparison between our approach and $BM25 + Sim_{d_{para}-Q}$ [32] on the TREC 2014 and 2015 CDS task

Method	infNDCG	infAP	MAP	R-Prec
2014 CDS task				
$BM25 + Sim_{d_{para}-Q}$	0.2618	0.0763	0.1579	0.1518
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2698	0.0935	0.1628	0.2067
	+ 3.06%	+ 22.54%*	+ 3.10%	+ 36.17%*
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.2748	0.0953	0.1645	0.2083
	+ 4.97%*	+ 24.90%*	+ 4.18%	+ 37.22%*
2015 CDS task				
$BM25 + Sim_{d_{para}-Q}$	0.2742	0.0657	0.1642	0.1491
$BM25 + SEM_{d_{para}-D_{PRF}^k}$	0.2980	0.0831	0.1758	0.2345
	+ 8.68%*	+ 26.48%*	+ 7.06%*	+ 57.28%*
$BM25 + SEM_{d_{sum}-D_{PRF}^k}$	0.2986	0.0842	0.1791	0.2408
	+ 8.90%*	+ 28.16%*	+ 9.07%*	+ 61.50%*

The results are obtained based on the *Summary* field

Table 6 Comparison to $SNUMedinfo$, the best automatic run in the TREC 2014 CDS task

Method	infNDCG	infAP
$SNUMedinfo$	0.2674	0.0659
$BM25 + SEM_{d-D_{PRF}^k}$	0.2830	0.0911

Results of $SNUMedinfo$ are taken from those reported in [6]. $BM25 + SEM_{d-D_{PRF}^k}$ is the best result of our approach on this dataset, as in Table 2. No statistical test is conducted due to unavailability of the per-query result of $SNUMedinfo$

remaining subset for testing. There is no overlap between the training and testing topics. Then the overall retrieval performance is obtained by averaging over the two test subsets of topics. Apart from the official TREC measure inferred NDCG (infNDCG) [27], we also report on other popular evaluation metrics in the CDS task, including Mean Average Precision (MAP) [7], R-Precision (R-Prec) [7] and inferred Average Precision (infAP) [27]. All statistical tests are based on the t test at the 0.05 significance level.

Table 7 Comparison to *WSU-IR*, the best automatic run in the TREC 2015 CDS task

Method	infNDCG	infAP
<i>WSU-IR</i>	0.2939	0.0842
<i>BM25</i> + <i>SEM</i> _{<i>d</i>-D_{PRF}^k}	0.3016 + 2.62%	0.0873 + 3.68%

Results of *WSU-IR* are taken from those reported in [2]. *BM25* + *SEM*_{*d*-D_{PRF}^k} is the best result of our approach on this dataset, as in Table 3. The difference between the two approaches is not statistically significant

5 Evaluation Results

In this section, we present the evaluation results of our proposed CDS method. Table 2 presents the evaluation results of the TREC 2014 CDS task using the *Summary* and *Description* fields, and Table 3 and present the evaluation results of the TREC 2015 CDS task A. Table 4 present the evaluation results of the TREC 2016 CDS task. Note that all the evaluation results are obtained by a twofold cross-validation based on the parity of the topic numbers. As described in Sect. 4.2, *BM25* + *SEM*_{*d*_{Para}-D_{PRF}^k} and *BM25* + *SEM*_{*d*_{Sum}-D_{PRF}^k} denote two different applications of our proposed CDS method, in which the embeddings of biomedical articles are generated by *Paragraph Embeddings* and *Term Summation*, respectively. Table 5 presents the comparison between our CDS method and the *BM25* + *Sim*_{*d*_{Para}-Q} method proposed in [32]. *BM25* is the baseline retrieval model used for verifying the effectiveness of our proposed feedback-based semantic relevance score. In addition, the comparisons between our approach and the best methods in the TREC 2014 and 2015 CDS tasks are presented in Tables 6 and 7, respectively. According to the results, we have the observations as follows.

First, our proposed feedback-based CDS method has statistically significant improvements over the baseline retrieval model *BM25* in most cases, which indicates the effectiveness of integrating semantic evidence into the frequency-based statistical models. Besides, according to Tables 6 and 7, our CDS method outperforms the best automatic methods in both TREC 2014 and 2015 CDS tasks. This observation is promising in that a simple linear interpolation of the classical *BM25* model and our proposed semantic relevance score could have scored the best run in those tasks.

Second, according to Table 5, our CDS method outperforms the method *BM25* + *Sim*_{*d*_{Para}-Q} proposed in [32], which integrates semantic evidence by measuring the cosine similarity between the embeddings of the patient record and biomedical article. As described in Sect. 1,

patient records are much shorter than full-text biomedical articles, such that patient record is a weak indicator of relevance, thus our feedback-based CDS method is expected to outperform the method *BM25* + *Sim*_{*d*_{Para}-Q}.

Third, comparing the two different ways of generating the article embeddings, *Term Summation* has a better performance than *Paragraph Embeddings* in most cases. As the full-text biomedical articles are usually very long, which contain large amount of irrelevant information, the mechanism of *Paragraph Embeddings* that considering the entire verbose texts while training embeddings may results in the sparse distribution of the semantic information in the embeddings of articles, such that the embeddings of articles generated by *Paragraph Embeddings* is not suitable to represent semantic relevance for long texts. In contrast, *Term Summation* generates embeddings of biomedical articles by only considering the top-*k* most informative words in the articles, which effectively reduces irrelevant information in the embeddings of biomedical articles.

Finally, comparing between the evaluation results obtained by using the *Summary* and *Description* fields in 2014 and 2015 CDS tasks, although using the *Description* field as queries obtained worse baseline retrieval results, the final performance of using *Description* field by integrating semantic evidence is better than using the *Summary* field in most cases. One possible reason is that the *Description* field is much longer than the *Summary* field, such that the relevant biomedical articles are returned by content-based retrieval models with relatively low ranking. By integrating the semantic evidence of relevance, the lowly ranked relevant documents are promoted in the ranking list which leads to improved retrieval performance. In 2016 CDS task, the final result by integrating semantic relevance in the *Description* field dose not outperform the result in the *Summary* field due to the extremely low baseline, but the improvement over the baseline method is statistically significant. The result in the *Note* field also shows the effectiveness of our proposed method. From our experience, the setting of parameter *k* (the number of top-*k* documents) has important impact on the effectiveness. In our experiments, this parameter is set by tuning on training set. According to the results obtained, it is suggested to set this parameter to 100, 10 and 80 on the 2014–2016 datasets, respectively.

5.1 Application of the Semantic Relevance Score to Other State-of-the-Art Methods

In this section, we use the best TREC run in 2015, *WSU-IR*, as the baseline to examine if our proposed method can still improve over the strongest baseline as far as we are aware of. We do not conduct the same comparison to

Table 8 The evaluation results on the TREC 2015 CDS task A

Method	infNDCG	infAP	MAP	R-Prec
Automatic run				
<i>wsuirdaa</i>	0.2939	0.0842	0.1864	0.2306
<i>wsuirdaa</i> + $SEM_{d_{para}-D_{PRF}^k}$	0.3130 + 6.50%*	0.0896 + 6.41%*	0.1905 + 2.20%	0.2396 + 3.90%
<i>wsuirdaa</i> + $SEM_{d_{sum}-D_{PRF}^k}$	0.3157 + 7.42%*	0.0898 + 6.65%*	0.1926 + 3.33%	0.2469 + 7.07%*
Manual run				
<i>wsuirdma</i>	0.3109	0.0880	0.1968	0.2493
<i>wsuirdma</i> + $SEM_{d_{para}-D_{PRF}^k}$	0.3265 + 5.02%*	0.0940 + 6.82%*	0.2015 + 2.39%	0.2605 + 4.49%
<i>wsuirdma</i> + $SEM_{d_{sum}-D_{PRF}^k}$	0.3335 + 7.27%*	0.0963 + 9.43%*	0.2054 + 4.37%	0.2643 + 6.02%*

The difference in percentage is measured against the baseline retrieval model *wsuirdaa* and *wsuirdma* [2]. A statistically significant difference is marked with a *. The best result of each evaluation metric is in bold

Table 9 The evaluation results on the TREC 2015 CDS task A

Method	infNDCG	infAP	MAP	R-Prec
Automatic run				
<i>wsuirdaa</i> + $SEM_{d_{LDA}-D_{PRF}^k}$	0.2963	0.0853	0.1864	0.2306
<i>wsuirdaa</i> + $SEM_{d_{para}-D_{PRF}^k}$	0.3130 + 5.64%*	0.0896 + 5.04%*	0.1905 + 2.20%	0.2396 + 3.90%
<i>wsuirdaa</i> + $SEM_{d_{sum}-D_{PRF}^k}$	0.3157 + 6.55%*	0.0898 + 5.28%*	0.1926 + 3.33%	0.2469 + 7.07%*
Manual run				
<i>wsuirdma</i> + $SEM_{d_{LDA}-D_{PRF}^k}$	0.3117	0.0887	0.1970	0.2494
<i>wsuirdma</i> + $SEM_{d_{para}-D_{PRF}^k}$	0.3265 + 4.75%*	0.0940 + 5.98%*	0.2015 + 2.28%	0.2605 + 4.45%
<i>wsuirdma</i> + $SEM_{d_{sum}-D_{PRF}^k}$	0.3335 + 6.99%*	0.0963 + 8.57%*	0.2054 + 4.26%	0.2643 + 5.97%*

The difference in percentage is measured against *wsuirdaa*(*wsuirdma*) + $SEM_{d_{LDA}-D_{PRF}^k}$. A statistically significant difference is marked with a *. The best result of each evaluation metric is in bold

Table 10 Statistics about the test collections and topics

Collection	TREC Task	Topics	# of Topics	# of Docs
disk1&2	1, 2, 3 ad hoc	51–200	150	741,856
disk4&5	Robust 2004	301–450 601–700	250	528,155
WT10G	9, 10 Web	451–550	100	1,692,096
GOV2	2004-2006	701–850	150	25,178,548
	Terabyte Ad-hoc			
ClueWeb09 B	2009-2011 Web	wt1–150	150	50,220,423

SNUMedinfo, the best TREC CDS run in 2014, due to unavailability of per-query results. In addition, the latent Dirichlet allocation (LDA) model [5] is applied to generate the distributed representations of biomedical articles for comparison with the neural embedding model Word2Vec in our study.

Table 8 presents the evaluation results based on the automatic and manual runs submitted by WSU-IR [2] in the TREC 2015 CDS Task A. *wsuirdaa* and *wsuirdma* in Table 8 are the submitted automatic and manual runs, respectively, and are used as our strong baselines. *wsuirdaa* + $SEM_{d_{para}-D_{PRF}^k}$ and *wsuirdaa* + $SEM_{d_{sum}-D_{PRF}^k}$

Table 11 Comparison to *BM25*

Model	disk1&2	disk4&5	WT10G	GOV2	CW09B
<i>BM25</i>	0.2408	0.2534	0.2123	0.3008	0.2251
<i>BM25</i> + $SEM_{d_{LDA}-D_{PRF}^k}$	0.2517 + 4.53%*	0.2675 + 5.56%*	0.2158 + 1.65%	0.3193 + 6.15%*	0.2306 + 2.44%*
<i>BM25</i> + $SEM_{d_{TFIDF}-D_{PRF}^k}$	0.2477 + 2.87%	0.2554 + 0.79%	0.2187 + 3.01%	0.3043 + 1.16%	0.2311 + 2.67%*
<i>BM25</i> + $SEM_{d_{Para}-D_{PRF}^k}$	0.2820 + 17.11%*	0.2862 + 12.94%*	0.2427 + 14.32%*	0.3138 + 4.32%*	0.2452 + 8.93%*
<i>BM25</i> + $SEM_{d_{Sum}-D_{PRF}^k}$	0.2727 + 13.25%*	0.2796 + 10.34%*	0.2423 + 14.13%*	0.3184 + 5.85%*	0.2404 + 6.80%*

The results on ClueWeb09B (CW09B) is reported in nDCG@20, and the rest are reported in MAP. A statistically significant difference is marked with a *. The best result on each collection is in bold

correspond to applying *Paragraph Embeddings* and *Term Summation*, respectively, the same to *wsuirdma* + $SEM_{d_{Para}-D_{PRF}^k}$ and *wsuirdma* + $SEM_{d_{Sum}-D_{PRF}^k}$. According to the results, we can see that there are still statistically significant improvements over the strong baselines *wsuirdaa* and *wsuirdma* in most cases when applying both *Paragraph Embeddings* and *Term Summation*, indicating the effectiveness of our proposed semantic relevance score.

Tables 9 presents the comparison between Word2Vec and LDA in generating the distributed representations of biomedical articles. The number of topics in LDA is set to 100 as used in [14]. *wsuirdaa* + $SEM_{d_{LDA}-D_{PRF}^k}$ and *wsuirdma* + $SEM_{d_{LDA}-D_{PRF}^k}$ in Table 9 corresponds to applying LDA model for generating the article representations, on top of the best TREC CDS runs in 2015. According to the results, there are statistically significant improvements over LDA in most cases when Word2Vec is utilized to generate the article embeddings. In fact, our experience indicates that the optimal value of hyper-parameter λ (See Equation (9)) when applying LDA model is usually 1, such that the semantic relevance score does not work when LDA is used. Therefore, we may conclude that

Word2Vec is more suitable than LDA for estimating the semantic similarity between biomedical articles.

6 Experimental Results on Other IR Test Collections

In addition to the CDS task, we further evaluate our proposed method on standard IR test collections in this section.

6.1 Experimental Settings

We use five standard TREC test collections in our experiments, and the basic statistics about the test collections and topics are given in Table 10. Documents are preprocessed by removing all HTML tags, standard English stopwords are removed and the test collections are stemmed using Porter's English stemmer. Each topic contains three fields, i.e., title, description and narrative, and we only use the title field. The title-only queries are very short which is usually regarded as a realistic snapshot of real user queries.

Table 12 Comparison to *BM25* with *Rocchio's PRF* ($BM25_{PRF}$)

Model	disk1&2	disk4&5	WT10G	GOV2	CW09B
<i>BM25_{PRF}</i>	0.3083	0.2966	0.2445	0.3430	0.2536
<i>BM25_{PRF}</i> + $SEM_{d_{LDA}-D_{PRF}^k}$	0.3084 + 0.03%	0.2966 + 0.0%	0.2446 + 0.04%	0.3479 + 1.43%	0.2596 + 2.37%*
<i>BM25_{PRF}</i> + $SEM_{d_{TFIDF}-D_{PRF}^k}$	0.3093 + 0.32%	0.2967 + 0.03%	0.2445 + 0.0%	0.3430 + 0.0%	0.2587 + 2.01%
<i>BM25_{PRF}</i> + $SEM_{d_{Para}-D_{PRF}^k}$	0.3110 + 0.88%	0.2990 + 0.81%	0.2541 + 3.93%*	0.3484 + 1.57%	0.2635 + 4.87%*
<i>BM25_{PRF}</i> + $SEM_{d_{Sum}-D_{PRF}^k}$	0.3105 + 0.71%	0.2985 + 0.64%	0.2541 + 3.93%*	0.3523 + 2.71%*	0.2677 + 5.42%*

The results on ClueWeb09B (CW09B) is reported in nDCG@20, and the rest are reported in MAP. A statistically significant difference is marked with a *. The best result on each collection is in bold

Table 13 The comparison between our approach and the locally trained method in [12] on $nDCG@10$

Model	disk1&2	disk4&5	CW09B
Locally-trained	0.563	0.517	0.258
Our method	0.5779	0.5261	0.2633

The results of the locally trained method are the best results reported in [12]. The results of our approach are obtained by $BM25 + SEM_{d_{para}-D_{PRF}^k}$

For each test collection, the Skip-gram model of Word2Vec or Para2Vec toolkit with negative sampling is utilized to generate word and document embeddings, which are trained by stochastic gradient ascent. The window size is set to 10 for Skip-gram model as recommended by [21]. The number of dimensions of the embeddings are set to 300. From our experience, with a wide range of possible settings, changing the number of dimensions of the word and document embeddings has little impact on the retrieval performance.

In our experiments, we evaluate our approach against $BM25_{PRF}$. In addition to the above baseline, the topic model LDA [5], and TF-IDF [18] are compared to Word2Vec or Para2Vec in generating the vector representations of documents in our experiments. The baseline models used in our experiments are optimized by grid search [4].

On each collection, we evaluate by a twofold cross-validation. The queries for each test collection are split into two equal-size subsets by parity in odd or even topic numbers. In each fold, one subset is used for training, and the other is used for test. The results reported in the paper are averaged over queries in the two test subsets. There is no overlap between the training and test subsets. We report on the official TREC evaluation metrics, including Mean Average Precision (MAP) [7] on disk1&2, disk4&5, WT10G, and GOV2, and $nDCG@20$ [7] on ClueWeb09 B. We use the official TREC evaluation metrics as we trust the TREC organizers to pick the appropriate measures for different retrieval tasks. All statistical tests are based on the t test at the 0.05 significance level.

6.2 Results

Table 11 presents the results against the classical $BM25$ model. According to the results, the integration of semantic relevance score (i.e., SEM) has statistically significant improvements over $BM25$ in all cases, indicating the effectiveness of our approach.

Table 12 presents the evaluation results against $BM25$ with Rocchio's PRF method. It is encouraging to see that statistically significant improvements are still observed with the use of PRF in most cases, especially on the three

Web collections, showing the effectiveness of our approach.

Table 11 also presents the comparison of three different models (i.e., Word2Vec or Para2Vec, LDA and TF-IDF) in generating the vector representations of documents. Out of the three models for document vector generating, Word2Vec or Para2Vec achieves the best effectiveness. LDA outperforms TF-IDF, but both of them are not as effective as Word2Vec or Para2Vec. The comparison result between Word2Vec or Para2Vec and LDA is consistent with the findings in other NLP tasks [11, 29]. As the TF-IDF vector representations of documents do not have the ability in capturing the semantic relation between texts, the comparison results between TF-IDF and the other two models can be expected.

We also compare the results of the proposed approach with a state-of-the-art query expansion approach based on locally trained embeddings [12]. This approach can possibly deal with the problem of multiple degrees of similarity by training the word embeddings on only the top-1000 documents. However, the on-line computational overhead could be an issue in practice since the word embeddings are trained on a per-query basis. As only $nDCG@10$ is used in [12], Table 13 compares the best $nDCG@10$ reported in [12] with our approach on each of the three publicly available TREC collections. From the comparison results we can see that our method consistently has better results over [12] on all three TREC test collections. Moreover, it can be observed that the Paragraph Embedding method outperforms Term Summation in most cases. A possible explanation is that the documents in traditional newswire or Web collections are in general much shorter and more coherent than the scientific articles. In this case, training a document embedding as a whole instead of summing up individual term embeddings may result in better document representations, and consequently better retrieval performance.

7 Conclusions and Future Work

In this paper, we have proposed a novel feedback-based CDS method, which integrates the semantic similarity between a biomedical article and the corresponding pseudo-relevance feedback set into frequency-based models. Experimental results show that integrating semantic evidence of relevance can indeed significantly improve the retrieval performance over the existing CDS approaches, including the best TREC results. In addition, a simple linear combination of the classical $BM25$ model with our proposed semantic relevance score ($BM25 + SEM_{d-D_{PRF}^k}$) would have achieved the best automatic runs on the TREC

2014 and 2015 CDS tasks. Compared to *Paragraph Embeddings*, *Term Summation* is more suitable to generate the embeddings of biomedical articles, due to the ability of reducing irrelevant information in the embeddings of biomedical articles. The comparison between Word2Vec and LDA shows that Word2Vec is more suitable than LDA for estimating the semantic similarity between biomedical articles. In future research, we plan to utilize the semantic relevance score for query expansion to further improve the performance of a CDS system.

Acknowledgements This work is supported by the National Natural Science Foundation of China (61472391). We would like to thank the authors of [2] for kindly sharing their TREC runs.

Compliance with Ethical Standards

Conflict of interest This is a statement that none of the authors have any Conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Abacha A, Khelifi S (2015) LIST at TREC 2015 clinical decision support track: question analysis and unsupervised result fusion. In: TREC
2. Balaneshinkordan S, Kotov A, Xisto R (2015) WSU-IR at TREC 2015 clinical decision support track: joint weighting of explicit and latent medical query concepts from diverse sources. In: TREC
3. Bengio Y, Schwenk H, Sencal J, Morin F, Gauvain J (2006) Neural probabilistic language models. Springer, Berlin, pp 137–186
4. Bergstra J, Bardenet R, Kgl B, Bengio Y (2011) Algorithms for hyper-parameter optimization. In: Advances in neural information processing systems, pp 2546–2554
5. Blei D, Ng A, Jordan M (2003) Latent dirichlet allocation. J Mach Learn Res 3(Jan):993–1022
6. Choi S, Choi J (2014) SNUMedinfo at TREC CDS track 2014: medical case-based retrieval task. Technical report, DTIC document
7. Chowdhury G (2007) TREC: experiment and evaluation in information retrieval. Online information review (5)
8. Collobert R, Weston J, Bottou L, Karlen M, Kavukcuoglu K, Kuksa P (2011) Natural language processing (almost) from scratch. J Mach Learn Res 12(Aug):2493–2537
9. Cummins R (2015) Clinical decision support with the SPUD language model. In: TREC
10. Cummins R, Paik J, Lv Y (2015) A plya urn document language model for improved information retrieval. ACM Trans Inf Syst (TOIS) 33(4):21
11. Dai A, Olah C, Le Q (2015) Document embedding with paragraph vectors. CoRR, abs/1507.07998
12. Diaz F, Mitra B, Craswell N (2016) Query expansion with locally-trained word embeddings. In: Proceedings of ACL, pp 1–11
13. Goldberg Y, Levy O (2014) word2vec explained: deriving Mikolov et al.s negative-sampling word-embedding method. CoRR, abs/1402.3722
14. Goodwin T, Harabagiu S (2014) UTD at TREC 2014: query expansion for clinical decision support. Technical report, DTIC document
15. Gurulingappa H, Toldo L, Schepers C, Bauer A, Megaro G (2016) Semi-supervised information retrieval system for clinical decision support. In: TREC
16. Hui K, He B, Luo T, Wang B (2011) A comparative study of pseudo relevance feedback for ad-hoc retrieval. In: Proceedings of ICTIR, pp 318–322
17. Le Q, Mikolov T (2014) Distributed representations of sentences and documents. CoRR, abs/1405.4053
18. Manning C, Raghavan P, Schtze H (2008) Introduction to information retrieval. Cambridge University Press, Cambridge
19. Metzler D, Croft W (2005) A Markov random field model for term dependencies. In: Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval, pp 472–479. ACM
20. Mikolov T, Yih W, Zweig G (2013) Linguistic regularities in continuous space word representations. In: HLT-NAACL
21. Mikolov T, Chen K, Corrado G, Dean J (2013) Efficient estimation of word representations in vector space. CoRR, abs/1301.3781
22. Mnih A, Hinton G (2008) A scalable hierarchical distributed language model. In: Conference on neural information processing systems. Vancouver, British Columbia, Canada, pp 1081–1088
23. Palotti J, Hanbury A (2015) TUW @ TREC clinical decision support track 2015. In: TREC
24. Ponte J, Croft W (1998) A language modeling approach to information retrieval. In: Proceedings of the 21st annual international ACM SIGIR conference on research and development in information retrieval, pp 275–281. ACM
25. Roberts K, Simpson M, Voorhees E, Hersh W (2015) Overview of the TREC 2015 clinical decision support track. In: TREC
26. Robertson S, Walker S, Beaulieu M, Gatford M, Payne A (1996) Okapi at TREC-4. TREC, pp 73–96
27. Simpson M, Voorhees E, Hersh W (2014) Overview of the TREC 2014 clinical decision support track. Technical report, DTIC document
28. Song Y, He Y, Hu Q, He L (2015) ECNU at 2015 CDS track: two re-ranking methods in medical information retrieval. In: TREC
29. Sun F, Guo J, Lan Y, Xu J, Cheng X (2016) Semantic regularities in document representations. CoRR abs/1603.07603
30. Turian J, Ratnoff L, Bengio Y (2010) Word representations: a simple and general method for semi-supervised learning. In: Proceedings of the 48th annual meeting of the association for computational linguistics. Association for Computational Linguistics, pp 384–394
31. Vuli I, Moens M (2015) Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In: The international ACM SIGIR conference, pp 363–372
32. Yang C, He B (2016) A novel semantics-based approach to medical literature search. In: 2016 IEEE international conference on bioinformatics and biomedicine (BIBM). IEEE, pp 1616–1623
33. Yang C, He B, Xu J (2017) Integrating feedback-based semantic evidence to enhance retrieval effectiveness for clinical decision support. In: Proceedings of APWEB-WAIM, pp 1–15