

Keyword Embeddings for Query Suggestion

Jorge Gabín^{1,2}[0000-0002-5494-0765], M. Eduardo Ares¹[0000-0003-3807-5692],
and Javier Parapar²[0000-0002-5997-8252]

¹ Linknovate Science, Rúa das Flores 33, Roxos, Santiago de Compostela, A Coruña,
15896, Spain

² IRLab, CITIC, Computer Science Department, University of A Coruña, A Coruña,
15071, Spain

Abstract. Nowadays, search engine users commonly rely on query suggestions to improve their initial inputs. Current systems are very good at recommending lexical adaptations or spelling corrections to users' queries. However, they often struggle to suggest semantically related keywords given a user's query. The construction of a detailed query is crucial in some tasks, such as legal retrieval or academic search. In these scenarios, keyword suggestion methods are critical to guide the user during the query formulation. This paper proposes two novel models for the keyword suggestion task trained on scientific literature. Our techniques adapt the architecture of **Word2Vec** and **FastText** to generate keyword embeddings by leveraging documents' keyword co-occurrence. Along with these models, we also present a specially tailored negative sampling approach that exploits how keywords appear in academic publications. We devise a ranking-based evaluation methodology following both known-item and ad-hoc search scenarios. Finally, we evaluate our proposals against the state-of-the-art word and sentence embedding models showing considerable improvements over the baselines for the tasks.

Keywords: Keyword suggestion · Keyword embeddings · Negative sampling · Academic search.

1 Introduction

The use of word embeddings [6] has improved the results of many Natural Language Processing (NLP) tasks, such as name entity recognition, speech processing, part-of-speech tagging, semantic role labelling, chunking, and syntactic parsing, among others. These techniques represent words as dense real-valued vectors which preserve semantic and syntactic similarities between words. More recently, the so-called document and sentence embedding models allow the computing of embeddings for larger pieces of text directly (instead of doing it word by word), getting state-of-the-art results on different tasks.

Nowadays, search engines are outstanding in recommending lexical adaptations or corrections to users' queries [14,17]. However, there is still room for improvement in suggesting semantically similar phrases to users' inputs. The few

systems that address this task do it by simply using existing word or sentence embedding models, which leads to poor keyword suggestions.

Keyword suggestion, which consists in recommending keywords similar to the user’s input, is critical in search scenarios where the completeness of the query clauses may dramatically affect the recall, such as academic or legal search [24,13]. An incomplete query may end up in a *null search session* [15], that is, the system presents an empty result list to the user. In tasks where use cases are recall-oriented, having an incomplete or even empty results set greatly diminishes the search experience.

Having good keyword suggestion models in a keyword-based search engine has many benefits. (i) It will help to identify the query intent [10], as users will be able to refine their search by adding new query clauses that may be clarifying. (ii) It will promote serendipity [5,1], as users may see a recommended keyword that they were not considering but perfectly fits their search interest. (iii) It will help prevent null sessions or incomplete results lists. (iv) Systems may use semantically similar keywords to the user input to perform query expansion [7,22] without further user interaction.

In this paper, we leverage `Word2Vec` [19] and `FastText` [4] models’ architecture to generate keyword embeddings instead of word embeddings, this meaning that embeddings represent sequences of words rather than a single word, similar to sentence embeddings. Unlike the base models, which use bag-of-words to build the representations, our approaches are based on bag-of-keywords. Thus, the proposed models exploit the annotated keywords’ co-occurrence in the scientific literature for learning the dense keyword representations.

Our main aim for producing keyword embeddings is to represent concepts or topics instead of representing just words’ semantics or even their contextual meaning like word embedding models do. The keywords’ conceptual or topical semantics rely on an explicit human annotation process. First, the annotator will only select proper or commonly used descriptors for the presented documents. In this way, we may assume that those keywords are a good proxy when searching for the documents’ topics. Second, authors use a limited number of keywords when annotating documents, so the strength of the semantic relationships among them is higher than the traditional approach of word co-occurrence in windows of free text.

Along with these models, we propose a new method to perform negative sampling, which leverages connected components to select the negative inputs used in the training phase. This method uses the keywords co-occurrence graph to extract the connected components and then select inputs that are not in the same connected component as negative samples.

To evaluate our models’ performance, we compare them against a set of baselines composed of state-of-the-art word and sentence embedding models. We both trained the baselines in the keyword data and used the pre-trained models. In particular, we further trained `Word2Vec` [19] and `Sentence-BERT` [21] (SBERT) on the task data, and we used SBERT, `SciBERT` [2] and `FastText` [4] base models. We carry out the evaluation in `Inspec` [11] and `KP20k` [18], two classical keyword

extraction datasets compound of scientific publications which count with a set of human-annotated keywords (for now on, we refer to keywords annotated by documents’ authors or professional annotators as annotated keywords). The results show significant improvements for the task over state-of-the-art word and sentence embedding models.

The main contributions of our work can be summarized as follows:

- **Keywords2Vec**, a keyword embedding model based on **Word2Vec**.
- **FastKeywords**, a keyword embedding model based on **FastText**.
- A new method for negative sampling based on connected components.
- Baselines and evaluation methods for the similar keyword suggestion task.

2 Related Work

This section briefly overviews the existing word and sentence embedding models, particularly those used as baselines in the evaluation stage, along with some previous work related to the keyword suggestion task.

Dense vector representations, a.k.a. embeddings, have an appealing, intuitive interpretation and can be the subject of valuable operations (e.g. addition, subtraction, distance measures, etc.). Because of those features, embeddings have massively replaced traditional representations in most Machine Learning algorithms and strategies. Many word embedding models, such as **Word2Vec** [19], **GloVe** [20] or **FastText** [4], have been integrated into widely used toolkits, resulting in even more precise and faster word representations.

Word2Vec [19] was one of the first widely used neural network-based techniques for word embeddings. These representations preserve semantic links between words and their contexts by using the surrounding words to the target one. The authors proposed two methods for computing word embeddings [19]: skip-gram (SG), which predicts context words given a target word, and continuous bag-of-words (CBOW), which predicts a target word using a bag-of-words context.

FastText [4] is a **Word2Vec** add-on that treats each word as a collection of character n -grams. **FastText** can estimate unusual and out-of-vocabulary words thanks to the sub-word representation. In [12], authors employed **FastText** word representation in conjunction with strategies such as bag of n -gram characteristics and demonstrated that **FastText** outperformed deep learning approaches while being faster.

Sentence embeddings surged as a natural progression of the word embedding problem. Significant progress has been made in sentence embeddings in recent years, particularly in developing universal sentence encoders capable of producing good results in a wide range of downstream applications.

Sentence-BERT [21] is one of the most popular sentence embedding models and state-of-the-art on the sentence representation task. It is a modification of the **BERT** [9] network using siamese and triplet networks that can derive semantically meaningful sentence embeddings.

There are other sentence embedding models based on pre-trained language models. An example of those language models which achieves excellent results when working with scientific data is `SciBERT`. `SciBERT` [2] is an adaptation of `BERT` [9] to address the lack of high-quality, large-scale labelled scientific data. This model leverages `BERT` unsupervised pre-training capabilities to further train the model on a large multi-domain corpus of scientific publications, producing significant improvements in downstream scientific NLP tasks.

Even though word and sentence embeddings have been widely studied, the work on keyword embeddings is still limited. Researchers have employed them mainly on tasks like keyword extraction, phrase similarity or paraphrase identification. Several approaches for the keyword extraction task, like `EmbedRank` [3], directly rely on pre-trained sentence embedding models to rank the extracted keywords. However, other approaches such as `Key2Vec` [16] train their own model for the phrase embedding generation. In particular, the authors propose directly training multi-word phrase embeddings using `FastText` instead of a classic approach that learns a model for unigram words combining the words' dense vectors to build multi-word embeddings later.

Yin and Schütze [26] presented an embedding model for generalized phrases to address the paraphrase identification task. This approach aims to train the `Word2Vec` SG model without any modification to learn phrase embeddings. They pre-process the corpus by reformatting the sentences with the continuity information of phrases. The final collection contains two-word phrases whose parts may occur next to each other (continuous) or separated from each other (discontinuous).

The phrase semantic similarity task is akin to the one we address in this paper. Many sentence embedding models, including `SBERT`, are pre-trained in that downstream task. In [27], the authors present a composition model for building phrase embeddings with semantic similarity in mind. This model, named `Feature-rich Compositional Transformation (FCT)`, learns transformations for composing phrase embeddings from the component words based on extracted features from a phrase.

Finally, the keyword suggestion task is also a sub-type of query suggestion where all the input and output queries are represented as keywords. Previous works in query term suggestion did not leverage the power of keyword co-occurrence to recommend new terms. Instead, existing query suggestion systems usually approach this task by suggesting terms extracted from the document's content without relying on the relations between these terms. For example, in [25], authors propose indexing a set of terms extracted from the corpus to rank them using a language model given an input query. The problem with these approaches is that they depend on the appearance of semantically related terms in the analyzed documents.


3 Proposal

In this section, we present two novel keyword embedding models that leverage `Word2Vec` [19] and `FastText` [4] architectures to produce keyword embeddings. We name these models `Keywords2Vec` and `FastKeywords`, respectively. Unlike `Word2Vec` and `FastText`, these models are not trained on rolling windows over the documents’ full text; instead, we only use combinations from the documents’ set of keywords as inputs.


The first of them, `Keywords2Vec`, modifies `Word2Vec` CBOB architecture to represent each keyword as one item. That is, we use keywords as token inputs instead of words. Additionally, we change how to perform the training of the models; we will explain it later for both proposals. We also evaluated the SG counterpart, but it performed considerably worse than the CBOB, so we do not report them here for brevity.

The second one, `FastKeywords`, adapts `FastText` CBOB variant in a more complex way. First, instead of working with words as the bigger information unit, it works with keywords. Second, it always selects each word of the keyword and the keyword itself as inputs, and then, during the n -grams selection process, it generates each word’s n -grams.


Taking the keyword “*search engine*” and $n = 3$ as an example, it will be represented by the following n -grams:



the special sequences for words:



and the special sequence for the whole keyword:



The model uses special sequences to capture the semantic meaning of both words and keywords. Finally, we implemented a weighting system to ignore “fill” n -grams used when a keyword does not have enough n -grams to fill the desired input size. The inclusion of this kind of n -grams is needed because the model requires the same input length on every iteration.

Another novel contribution of this paper is how we generate the training inputs. The model needs both positive and negative contexts for the target keyword. They are called positive and negative samples. For producing the positive samples, we select combinations of annotated keywords from the document to which the target keyword belongs. In the case of the negative samples, we represent the keywords’ co-occurrences in the dataset as a graph. In this graph, each keyword is a node and edges are created when two keywords appear together in a document. We select the negative samples from connected components different from that of the target keyword.

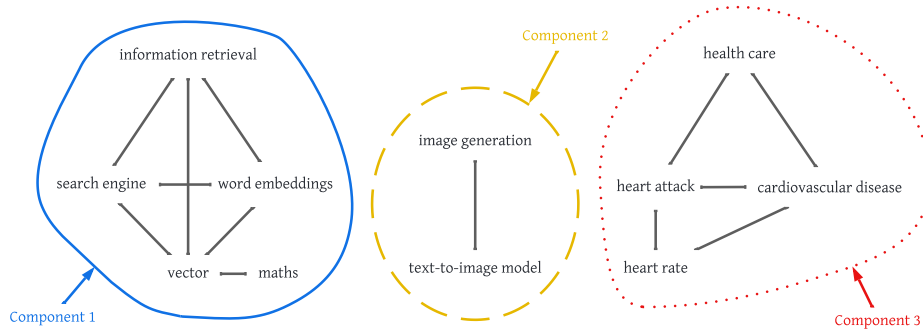


Fig. 1. Example of a keyword co-occurrence graph with its connected components.

Figure 1 shows an example of a keyword co-occurrence graph and its connected components. In that example, the positive samples for the keyword “*information retrieval*” would be combinations of “*search engine*”, “*word embeddings*” and “*vector*”, and its negative samples will always be extracted from connected components 2 or 3. Note that this means that the keyword “*maths*” will never be a negative sample for “*information retrieval*”, even though it does not co-occur with it in any document.

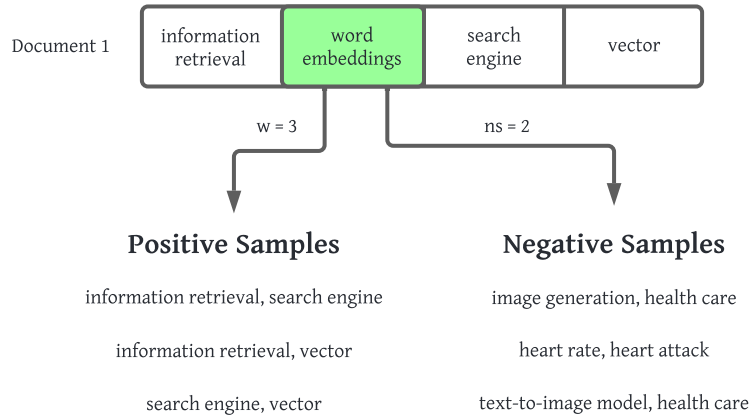


Fig. 2. Positive and negative samples generation example. (*The target keyword is highlighted*).

Figure 2 shows how the positive and negative sampling is performed for a document extracted from the graph shown in the Figure 1. Having the keyword “*word embeddings*” as the target, we first select its positive samples. In this example, we select every combination of 2 ($w - 1 = 2$) keywords that belong to the same document as the target. Then, for each pair of keywords we have

to select the two negative samples ($ns = 2$). Given the connected components shown in the Figure 1, we select two keywords that do not belong to the same connected component than the target keyword.

Formally, to train these models, we use a dataset of scientific documents (D). Each document ($d \in D$) contains a set of annotated keywords (K_d). Then, given a combination size ($w - 1$), which plays an analogous role to the window size in the original models, we compute the set of positive samples (K_d^{ps}) of a document d as follows:

$$K_d^{ps} = \left\{ (k_i, K_d^{psj}) \mid k_i \in K_d, K_d^{psj} \in \binom{K_d - \{k_i\}}{w - 1} \right\},$$

that is, for each keyword (k_i) of the document’s keywords set (K_d), we compute its positive samples (K_d^{ps}), which are the combinations of size $w - 1$ of the document’s keywords set excluding the target keyword ($K_d - \{k_i\}$). Finally, for each positive samples set (K_d^{psj}) we obtain a pair (k_i, K_d^{psj}).

As for the negative samples for each document (K_d^{ns}) we have followed the subsequent novel approach. First, we build the aforementioned keywords co-occurrence graph for the collection and compute its connected components. Then, for each pair (k_i, K_d^{psj}) we select as negative samples ns keywords belonging to a different connected component than the target keyword (k_i).

To adapt the previous process to our **FastKeywords** model, we have to generate n -grams for each context keyword (positive samples and negative samples) following the strategy explained before.

The **FastKeywords** model has two main advantages over **Keywords2Vec** because of the use of subword information:

- It will perform significantly better on large collections.
- It will be able to generate embeddings for keywords that are not in the training corpus.

Figure 3 shows the **FastKeywords** model’s architecture. As we may see, it follows the classic CBOW strategy where several context samples are fed to the model in order to predict the target keyword.

4 Experimental Setup

This section describes the datasets used during the training and evaluation of the models, the baselines used to compare our model, the evaluation process and metrics, and finally, the parameters and setup used to train our models.

Our objective is to suggest similar keywords to the user input in the search process. With this in mind, we designed two experiments that approach the keyword suggestion problem as a keyword ranking task. In particular, the evaluation considers if the model can find keywords that belong to the same document as the target keyword. The rationale for this evaluation is that keywords with which the authors annotate a document tend to be semantically related. Alternatively,

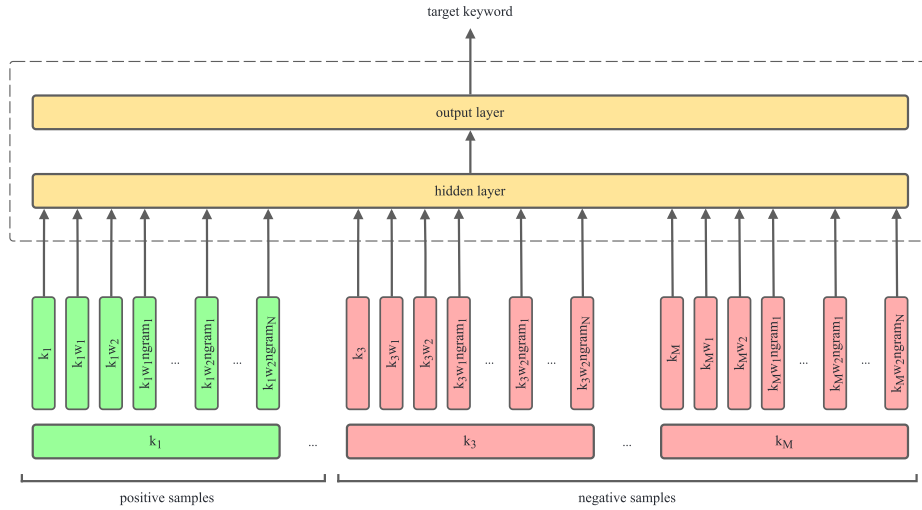


Fig. 3. FastKeywords model high-level architecture diagram.

we may perform user studies to evaluate the perceived quality of the suggestions. We leave that for future work.

4.1 Datasets

We selected two datasets commonly used on the classical keyword extraction task: the Inspec [11] and the KP20k [18] collections.

Inspec This collection consists of 2,000 titles and abstracts from scientific journal papers. The annotators assigned two sets of keywords to each document in this collection: controlled keywords that occur in the Inspec thesaurus and uncontrolled keywords that can be any suitable terms. We will only use the uncontrolled keywords created by professional indexers in our experiments, following the standard approach in the keyword extraction task.

KP20k This dataset is a keyword generation collection that includes 567,830 articles of the computer science domain obtained from various online digital libraries, including ACM Digital Library, ScienceDirect, Wiley and Web of Science. In this case, the authors made no differentiation during the annotation process, so we used all of the keywords in both training and testing.

4.2 Baselines

We include five baselines in our evaluation to compare our embedding models against other models that have already proven their effectiveness in capturing words or sentence syntax and semantic information.

The first baseline is the `Word2Vec` model trained in the Google News dataset and further tuned on each dataset to learn their vocabulary. That process was performed using the keywords set as a sentence so that the model could learn the contextual relationships between them.

The second baseline is the pre-trained SBERT *all-mpnet-base-v2*³ model fine-tuned in the evaluation datasets for the sentence similarity task. To fine-tune the model, we served as inputs pairs of keywords and a binary value of similarity (1 if keywords belong to the same document and 0 if they do not).

The other three baselines were not trained for the task. These baselines are:

- The pre-trained SBERT model, *all-mpnet-base-v2*.
- The pre-trained SciBERT model that uses the uncased version of the vocabulary.
- The English FastText model trained on Common Crawl⁴ and Wikipedia⁵.

4.3 Evaluation

As we mentioned before, to assess our models’ performance in the keyword suggestion task, we devise two evaluation strategies representing two specific cases in the academic search field.

The first task consists of retrieving all the keywords with which a document was annotated, given one of them. The keyword retrieval is done via cosine similarity between the query keyword and all the other keywords. Then, we compute each average precision and aggregate all the results to compute the MAP (Mean Average Precision). Specifically, we calculate MAP@20.

The second task we propose follows a masking problem. For each document, we mask one keyword from the keywords set and then use the remaining keywords to try to find the masked one. For each of the remaining keywords, we retrieve the most similar ones using cosine similarity. Then, combining the scores of all the non-masked keywords, we compute the final ranking. Finally, we compute each masked keyword’s reciprocal rank to calculate the MRR (Mean Reciprocal Rank). Precisely, we compute MRR@100. Note that for this evaluation method, we only used 50 random documents from the test split of each dataset.

Regarding methodologies, we evaluated both tasks under all-items and test-items approaches. In the all-items fashion, the keywords to be ranked are the whole set of keywords in the dataset, while in the test-items indexing approach, only the keywords in the test subset are ranked.

We also perform statistical significance tests for all the evaluations. On the one hand, for comparing more than two systems, we use the Randomised Tukey HSD Test [8,23]. On the other hand, we follow the classic permutation test approach when comparing a pair of systems. These tests verify that the family-wise error does not exceed the confidence level α .

³ <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

⁴ <https://commoncrawl.org/>

⁵ <https://dumps.wikimedia.org/>

Table 1. Models’ training parameters on the Inspec and KP20k datasets.

Model	Inspec					KP20k				
	Batch	Epochs	Dim.	w	ns	Batch	Epochs	Dim.	w	ns
Word2Vec	-	<i>early stopping</i>	300	6	-	-	20	300	6	-
SBERT	2^9	<i>early stopping</i>	768	-	4	2^9	20	768	-	4
Keywords2Vec	2^{17}	<i>early stopping</i>	300	3	4	2^{17}	20	300	3	4
FastKeywords	2^{15}	<i>early stopping</i>	300	3	4	2^{15}	20	300	3	4

4.4 Experimental Settings

This section will describe the parameters used to train the baselines and our proposed models. Table 1 shows a summary of the models’ training parameters detailed hereunder.

To train the `Word2Vec` model, we followed an early-stopping strategy for the Inspec dataset, while for KP20k, we trained the model for 20 epochs. For both datasets, the selected embedding size was 300. Finally, we used a window size of 6 and 1 as the minimum count (lowest frequency of a word to be considered).

On the other hand, as we mentioned before, we fine-tuned the `SBERT` model in the sentence similarity task. For each positive sample (two keywords that co-occur in a document), we selected four negative samples (two keywords that never co-occur in a document). Again, we followed an early-stopping strategy for the Inspec dataset, while for KP20k, we trained the model for 20 epochs. For both datasets, we used a batch size of 2^9 and an embedding size of 768.

We trained both our models using an early-stopping strategy for the Inspec dataset. On the other hand, when using the KP20k dataset, we trained each model for 20 epochs. Regarding the batch size, the `Keywords2Vec` model was trained with a batch size of 2^{17} on both datasets, while for the `FastKeywords` model, we used a batch size of 2^{15} on Inspec and KP20k. Also, both were trained to generate embeddings of size 300. The combination size and the number of selected negative samples were the same for both models on both datasets, 2 ($w = 3$) and 4, respectively. Finally, regarding `FastKeywords` parameters, we used a minimum n -gram size of 3, a maximum n -gram size of 6 and a maximum number of n -grams of 20.

When trained, we used the whole set of document keywords from the dataset in the tuning process of the models⁶. When not trained, we used the default pre-trained models for the baselines. In the evaluation process, we selected a subset of documents on which we performed the aforementioned tasks.

⁶ All experiments were run on an NVIDIA A100 GPU with 80GB of memory.

Table 2. MAP@20 for the document’s keywords identification task. Statistically significant improvements according to the Randomized Tukey HSD test ($permutations = 100,000; \alpha = 0.05$) are superscripted.

Model	Inspec		KP20K	
	All items	Test items	All items	Test items
SBERT (<i>no train</i>) (†)	0.0323	0.0544 [Ⓜ]	0.0039	0.0134 [Ⓜ]
SciBERT (<i>no train</i>) (‡)	0.0205	0.0345	0.0040	0.0110
FastText (<i>no train</i>) (Ⓜ)	0.0147	0.0217	0.0039	0.0095
Word2Vec (⌘)	0.0463 [Ⓜ]	0.0832 ^{†‡Ⓜ}	0.0048	0.0130 ^{‡Ⓜ}
SBERT (±)	0.1481 ^{†‡Ⓜ⌘}	0.5880 ^{†‡Ⓜ⌘}	0.0072 ^{†‡Ⓜ⌘}	0.0182 ^{†‡Ⓜ⌘}
Keywords2Vec (⊗)	0.8634 ^{†‡Ⓜ⌘±}	0.9090 ^{†‡Ⓜ⌘±}	0.0690 ^{†‡Ⓜ⌘±}	0.0918 ^{†‡Ⓜ⌘±}
FastKeywords (⊙)	0.8659 ^{†‡Ⓜ⌘±}	0.9161 ^{†‡Ⓜ⌘±}	0.0762 ^{†‡Ⓜ⌘±⊗}	0.1060 ^{†‡Ⓜ⌘±⊗}

5 Results

This section reports how **Keywords2Vec** and **FastKeywords** perform against the selected baselines and how they perform against each other. As we mentioned in Section 4.3 we report results using two tasks and two evaluation techniques.

Table 2 shows the results for the document’s keywords identification task. We can see that our models significantly outperform the established baselines in both datasets and evaluation approaches. This is especially remarkable in the case of the state-of-the-art sentence BERT model, which, even for the fine-tuned scenario and using embeddings with twice the dimensions, lies quite behind our proposals. As expected, using the test-items strategy produces better results since the keyword set where the search is performed is much smaller.

Moreover, we can see that the dataset size is a determinant factor in this task. A much larger set of keywords (this is the case of the KP20k dataset versus the Inspec dataset) greatly impacts the final score. This makes sense because increasing the keyword set size makes finding the ones we are looking for more challenging (something analogous happens to a lesser degree between the all-items and test-items results).

In terms of comparing **FastKeywords** against **Keywords2Vec** we can see that the former performs better and that the more considerable difference between them appears when the keyword set is the biggest (KP20k dataset). The main reason behind this relies on the capability of **FastKeywords** to leverage subword information to build the keywords embeddings.

Table 3 shows results for the masked keyword discovering task. These results confirm what we had already seen in the first one: our models significantly outperform all the established baselines in every experiment we performed. Again, we can see that increasing the keyword set length produces worse results as the task becomes more and more challenging. Again, when comparing the proposed models against each other on this task, we can see that **FastKeywords** performs

Table 3. MRR@100 for the masked keyword discovering task. Statistically significant improvements according to the Randomized Tukey HSD test ($permutations = 1,000,000$; $\alpha = 0.05$) are superscripted.

Model	Inspec		KP20K	
	All items	Test items	All items	Test items
SBERT <i>no train</i> (†)	0.0342	0.0479	0.0018	0.0152
SciBERT <i>no train</i> (‡)	0.0220	0.0345	0.0056	0.0129
FastText <i>no train</i> (⊕)	0.0068	0.0097	0.0023	0.0047
Word2Vec (‡)	0.0751	0.1187 [⊕]	0.0008	0.0079
SBERT (±)	0.1688 ^{†‡⊕}	0.5890 ^{†‡⊕‡}	0.0138	0.0448
Keywords2Vec (⊗)	0.8914 ^{†‡⊕‡±}	0.9100 ^{†‡⊕‡±}	0.0778 ^{†‡⊕‡±}	0.0828 ^{†‡⊕‡}
FastKeywords (⊙)	0.8988 ^{†‡⊕‡±}	0.9102 ^{†‡⊕‡±}	0.1402 ^{†‡⊕‡±⊗}	0.1467 ^{†‡⊕‡±⊗}

Table 4. Negative sampling methods comparison on the Inspec dataset. Statistically significant improvements according to the permutation test ($permutations = 1,000,000$; $\alpha = 0.05$) are superscripted.

Model	MAP@20		MRR@100	
	All items	Test items	All items	Test items
FastKeywords <i>random negative sampling</i> (⊕)	0.8420	0.9093	0.8750	0.8890
FastKeywords (⊙)	0.8659 [⊕]	0.9161 [⊕]	0.8988	0.9102 [⊕]

better than **Keywords2Vec** when increasing the dataset size, getting statistically significant improvements over the **Word2Vec**-based method.

Table 4 shows the performance of the proposed negative sampling method, comparing it with a **FastKeywords** model that uses a random negative sampling strategy. The results ratify that the proposed method works significantly better than a naive random approach in all cases except on MRR@100 using the all-items indexing strategy. For this case, the p -value is 0.06.

Finally, for illustrative purposes, Table 5 shows the top 10 nearest neighbours for the keyword “*information retrieval*” on the KP20k collection retrieved by both proposed models. We also show each keyword’s associated score, which represents the similarity with the query keyword.

6 Conclusions

This paper explored the potential of keyword embedding models in the keyword suggestion task. We also propose several baselines and new evaluation methods to assess the performance of the models, as not much previous work has been published for this task.

The proposed models adapt **Word2Vec** and **FastText** CBOW architectures to compute keyword embeddings instead of word embeddings. Along with these two

Table 5. Top 10 nearest neighbours for the keyword “*information retrieval*” on the KP20k collection (*test-items index*).

FastKeywords		Keywords2Vec	
keyword	score	keyword	score
ranking	0.9133	retrieval model	0.6045
query expansion	0.9104	collection selection	0.5587
text classification	0.9067	link topic detection	0.5472
relevance	0.9044	dempstershafer evidence theory	0.5418
text mining	0.9018	retrospective evidence event detection	0.5371
information extraction	0.9012	instance based learning	0.5361
relevance feedback	0.8980	query expansion	0.5282
knowledge discovery	0.8979	terminology extraction	0.5260
document clustering	0.8971	viral marketing	0.5253
text categorization	0.8970	query formulation	0.5248

keyword embedding models, we present a novel strategy for the negative sampling task, which leverages the potential of the keyword co-occurrence graph’s connected components to perform a better selection of the negative samples.

Results show that our methods significantly outperform the selected baselines on both evaluation datasets. We also demonstrated the potential of sub-keyword and sub-word information to represent keywords as embeddings. In future work we aim to:

- Use the designed weights system to give more relevance to full keywords and words than to n -grams.
- Assess the models’ performance using popularity-based negative sampling.
- Combine negative samples extracted from the target keyword connected component and from different connected components.
- Use special delimiters to differentiate if a word is a part of a keyword or a keyword itself or if a n -gram is a part of a word or a word itself.
- Train and test the models on non-scientific keyword-style annotated data.
- Study how the offline findings of this work align with live user testing.

Acknowledgements This work was supported by projects PLEC2021-007662 (MCIN/AEI/10.13039/501100011033, Ministerio de Ciencia e Innovación, Agencia Estatal de Investigación, Plan de Recuperación, Transformación y Resiliencia, Unión Europea-Next Generation EU) and RTI2018-093336-B-C22 (Ministerio de Ciencia e Innovación, Agencia Estatal de Investigación). The first and third authors also thank the financial support supplied by the Consellería de Cultura, Educación e Universidade Consellería de Cultura, Educación, Formación Profesional e Universidades (accreditation 2019-2022 ED431G/01, ED431B 2022/33) and the European Regional Development Fund, which acknowledges the CITIC Research Center in ICT of the University of A Coruña as a Research Center of the Galician University System. The fist author also acknowledges the support of grant DIN2020-011582 financed by the MCIN/AEI/10.13039/501100011033.

References

1. André, P., m. c. schraefel, Teevan, J., Dumais, S.T.: Discovery is never by chance: designing for (un)serendipity. In: Bryan-Kinns, N., Gross, M.D., Johnson, H., Ox, J., Wakkary, R. (eds.) Proceedings of the 7th Conference on Creativity & Cognition, Berkeley, California, USA, October 26-30, 2009. pp. 305–314. ACM (2009). <https://doi.org/10.1145/1640233.1640279>
2. Beltagy, I., Lo, K., Cohan, A.: SciBERT: A pretrained language model for scientific text. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). pp. 3615–3620. Association for Computational Linguistics, Hong Kong, China (Nov 2019). <https://doi.org/10.18653/v1/D19-1371>
3. Bennani-Smires, K., Musat, C., Hossmann, A., Baeriswyl, M., Jaggi, M.: Simple unsupervised keyphrase extraction using sentence embeddings. In: Korhonen, A., Titov, I. (eds.) Proceedings of the 22nd Conference on Computational Natural Language Learning, CoNLL 2018, Brussels, Belgium, October 31 - November 1, 2018. pp. 221–229. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/k18-1022>
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. Transactions of the association for computational linguistics **5**, 135–146 (2017)
5. Buchanan, S.A., Sauer, S., Quan-Haase, A., Agarwal, N.K., Erdelez, S.: Amplifying chance for positive action and serendipity by design. Proceedings of the Association for Information Science and Technology **57**(1), e288 (2020)
6. Camacho-Collados, J., Pilehvar, M.T.: From word to sense embeddings: A survey on vector representations of meaning. Journal of Artificial Intelligence Research **63**, 743–788 (2018)
7. Carpineto, C., Romano, G.: A survey of automatic query expansion in information retrieval. Acm Computing Surveys (CSUR) **44**(1), 1–50 (2012)
8. Carterette, B.A.: Multiple testing in statistical analysis of systems-based information retrieval experiments. ACM Trans. Inf. Syst. **30**(1) (mar 2012). <https://doi.org/10.1145/2094072.2094076>
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1423>
10. Hu, J., Wang, G., Lochovsky, F., Sun, J.t., Chen, Z.: Understanding user’s query intent with wikipedia. In: Proceedings of the 18th International Conference on World Wide Web. p. 471–480. WWW ’09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1526709.1526773>
11. Hulth, A.: Improved automatic keyword extraction given more linguistic knowledge. In: Proceedings of the 2003 conference on Empirical methods in natural language processing. pp. 216–223 (2003)
12. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. In: Lapata, M., Blunsom, P., Koller, A. (eds.) Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers. pp. 427–431. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/e17-2068>

13. Kim, Y., Seo, J., Croft, W.B.: Automatic boolean query suggestion for professional search. In: Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 825–834. SIGIR '11, Association for Computing Machinery, New York, NY, USA (2011). <https://doi.org/10.1145/2009916.2010026>
14. Li, M., Zhu, M., Zhang, Y., Zhou, M.: Exploring distributional similarity based models for query spelling correction. In: Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics. pp. 1025–1032 (2006)
15. Li, X., Schijvenaars, B.J.A., de Rijke, M.: Investigating queries and search failures in academic search. *Inf. Process. Manag.* **53**(3), 666–683 (2017). <https://doi.org/10.1016/j.ipm.2017.01.005>
16. Mahata, D., Kuriakose, J., Shah, R., Zimmermann, R.: Key2vec: Automatic ranked keyphrase extraction from scientific articles using phrase embeddings. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers). pp. 634–639 (2018)
17. Martins, B., Silva, M.J.: Spelling correction for search engine queries. In: International Conference on Natural Language Processing (in Spain). pp. 372–383. Springer (2004)
18. Meng, R., Zhao, S., Han, S., He, D., Brusilovsky, P., Chi, Y.: Deep keyphrase generation. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 582–592. Association for Computational Linguistics, Vancouver, Canada (Jul 2017). <https://doi.org/10.18653/v1/P17-1054>
19. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: Bengio, Y., LeCun, Y. (eds.) 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings (2013), <http://arxiv.org/abs/1301.3781>
20. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
21. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics (11 2019), <https://arxiv.org/abs/1908.10084>
22. Russell-Rose, T., Gooch, P., Kruschwitz, U.: Interactive query expansion for professional search applications. *Business Information Review* **38**(3), 127–137 (2021)
23. Sakai, T.: Laboratory experiments in information retrieval. *The information retrieval series* **40** (2018)
24. Verberne, S., Sappelli, M., Kraaij, W.: Query term suggestion in academic search. In: European Conference on Information Retrieval. pp. 560–566. Springer (2014)
25. Verberne, S., Sappelli, M., Kraaij, W.: Query term suggestion in academic search. In: de Rijke, M., Kenter, T., de Vries, A.P., Zhai, C., de Jong, F., Radinsky, K., Hofmann, K. (eds.) *Advances in Information Retrieval - 36th European Conference on IR Research, ECIR 2014, Amsterdam, The Netherlands, April 13-16, 2014. Proceedings. Lecture Notes in Computer Science*, vol. 8416, pp. 560–566. Springer (2014). https://doi.org/10.1007/978-3-319-06028-6_57
26. Yin, W., Schütze, H.: An exploration of embeddings for generalized phrases. In: Proceedings of the 52nd Annual Meeting of the Association for Computational

- Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Student Research Workshop. pp. 41–47. The Association for Computer Linguistics (2014). <https://doi.org/10.3115/v1/p14-3006>
27. Yu, M., Dredze, M.: Learning composition models for phrase embeddings. *Trans. Assoc. Comput. Linguistics* **3**, 227–242 (2015). https://doi.org/10.1162/tacl_a_00135