



Universiteit
Leiden
The Netherlands

Relevance feedback strategies for reducing review effort in recall-oriented neural information retrieval

Kats, T.; Putten, P.W.H. van der; Scholtes, J.; Oliehoek, F.A.; Kok, M.; Verwer, S.

Citation

Kats, T., Putten, P. W. H. van der, & Scholtes, J. (2024). Relevance feedback strategies for reducing review effort in recall-oriented neural information retrieval. *Communications In Computer And Information Science*, 22-39. doi:10.1007/978-3-031-74650-5_2

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/4296315>

Note: To cite this publication please use the final published version (if applicable).



Relevance Feedback Strategies for Reducing Review Effort in Recall-Oriented Neural Information Retrieval

Timo Kats¹, Peter van der Putten¹(✉), and Jan Scholtes²

¹ LIACS, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands
{t.p.a.kats,p.w.h.van.der.putten}@liacs.leidenuniv.nl

² Maastricht University, Minderbroedersberg 4-6, 6211 LK Maastricht, The Netherlands
scholtes@msc.nl

Abstract. In a number of information retrieval applications, such as patent search, literature review, and due diligence, preventing false negatives is more important than preventing false positives. However, approaches designed to reduce review effort, such as ‘technology assisted review’, can create false negatives, since these are often based on active learning systems that exclude documents automatically based on user feedback. To address this issue, we propose a recall-oriented approach to reducing review effort, through iteratively re-ranking the relevance rankings based on user feedback. We propose and experiment with various relevance feedback strategies. Best results are obtained by using a BERT-based dense-vector search for relevance rankings, and basing relevance feedback on cumulatively summing the queried and selected embeddings. Our results show that this method can reduce review effort between 17.85% and 59.04%, compared to a baseline approach of no feedback, given a fixed recall target.

Keywords: high recall information retrieval · dense-vector search · relevance feedback · neural information retrieval

1 Introduction

In a diverse range of information retrieval applications, reaching certain minimal recall goals is key, when identifying which documents in a large corpus are relevant to an information need. This is particularly important in contexts such as intelligence and law enforcement, searching for adverse effects of medicines, patent due diligence, or legal search applications. In these scenarios, preventing false negatives often takes precedence over preventing false positives [23].

Manually reviewing all documents is cost intensive and error-prone [17]. Therefore, technology is often used to reduce review effort. However, technolo-

gies aimed at reducing review effort, such as ‘technology assisted review’ can create false negatives, since they often rely on active learning systems that exclude documents automatically based on user feedback [10].

Therefore, in this research we propose a more recall-oriented approach to reducing review effort while cumulatively identifying relevant documents. For this, we combine two major components: textual similarity and relevance feedback. Textual similarity is used to produce relevance rankings, while relevance feedback iteratively refines these relevance rankings based on user feedback. Hence, our first research question are: ‘what text similarity methods are most suitable for relevance feedback?’, and building upon this: ‘to what extent can relevance feedback contribute to reducing review effort?’

For textual similarity, we evaluate TF-IDF and BERT-based dense-vector representations. For relevance feedback, we compare text-based and vector-based feedback strategies. Additionally, we experiment with varying levels of textual granularity in both components. In the text similarity component this is done through paragraph-based document rankings and in the relevance feedback component this is done through feedback amplification.

This combination of representation methods, feedback strategies and levels of granularity has elements that contribute to scientific novelty, and leads to substantial improvements in results. In fact, given a recall target of 80%, our proposed method reduces review effort between 17.85% and 59.04% compared to our baseline approach of no relevance feedback.

The remainder of this paper is structured as follows. Section 2 discusses the background. Next, we describe our methodology (Sect. 3) and experimental setup (Sect. 4). In Sect. 5 we share the results of our experiments, and the paper ends with a discussion section (Sect. 6) and conclusion (Sect. 7).

2 Background

In this section we discuss work related to the two main components of our method: text similarity methods and relevance feedback strategies. For the text similarity methods we discuss term-based and context-based approaches. Next, for the relevance feedback strategies, we review text-based and vector-based strategies.

2.1 Term-Based Similarity Methods

Term-based similarity methods compute text similarity through taking the common features between two pieces of text into account [6]. As a result, these methods do not incorporate contextual language properties like homonymy or synonymy in their similarity computations.

A simple and commonly used [12] implementation of term-based similarity is Jaccard similarity [11], which computes text similarity based on how many features two texts have in common divided by the total number of features across

both texts. The formula for this is given below. Here, A and B refer to the set of unique words for both documents.

$$Jaccard(A, B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B|} \quad (1)$$

However, a shortcoming of Jaccard similarity is that frequent terms (that are more likely to match) within a document are unlikely to be distinctive or important [5]. To deal with this shortcoming, there are two approaches. The first approach is based on manually filtering out frequent words based on a predefined list of words that are known to be common in a given language. These words are often referred to as ‘stop words’. Removing stop words when searching for textually similar documents generally has a beneficial effect [5], though this is task dependent [25]. Hence, we filter out stop words in our experiment.

The second approach is based on reducing the influence of words that are common in a corpus [5]. An advantage of this method compared to filtering stop words is that it is more dynamic. Certain words that are common within a specific context (e.g., the word ‘patient’ in medical data) might not be included in predefined lists of stop words. This approach is implemented in our TF-IDF-based similarity method through *inverse document frequency*. Terms are given a measure of uniqueness by dividing the number of documents in total by the number of documents that have a specific term. The formula for this metric is given below. Here, D refers to the number of documents in the dataset whereas d refers to the number of documents that contain term t . As a result, terms that appear in many documents (and are therefore less unique) are given a diminished value [22].

$$idf(t, D) = \log \frac{|D|}{1 + |\{d \in D : t \in d\}|} \quad (2)$$

2.2 Context-Based Similarity Methods

In this research we also use context-based similarity methods. In contrast to the term-based similarity methods mentioned previously, these methods do incorporate a form of semantics. This is based on the ‘distributional hypothesis’ [9], which is built on the idea that ‘words that occur in the same contexts tend to have similar meanings’ [6]. Given this hypothesis, this research uses pre-trained word embeddings that provide vector-based representations of texts. This enables us to compute the textual similarity based on the similarity between the vectors (e.g., with cosine similarity) [6].

In this category of pre-trained word embeddings, we specifically focus on BERT [8, 26]. The transformer-based architecture of BERT allows it to capture context-sensitive word properties (like homonymy and synonymy) in its embeddings. There are versions of BERT made for specific tasks. For example, Sentence-BERT (SBERT) is a BERT model specifically made for measuring text similarity using cosine distance [19]. The key advantage of SBERT over other BERT models for text similarity tasks is its reduction in computational

overhead. The researchers found that for finding the most similar pairs in a collection of 10,000 sentences, BERT would take ≈ 65 h whereas SBERT would take 5 s. Moreover, it enables a similarity search between larger bodies of text (like sentences and paragraphs) through mean-pooling the word embeddings.




Fig. 1. Simplified architecture of SBERT [19].

A simplified overview of this architecture can be found in Fig. 1. Here, the (word-level) BERT embeddings are first mean-pooled to create paragraph or sentence level embeddings. Thereafter, these embeddings are compared through their cosine similarity. Note, for the mean-pooling process, text that exceeds 384 words in length is truncated [19]. As a result, this approach cannot be implemented on a level of text granularity that exceeds this amount.

2.3 Text-Based Feedback Strategies

Relevance feedback refers to changing the relevance ranking through user feedback, generally in multiple iterations [16]. A commonly used text-based relevance feedback strategy is ‘keyword expansion’. Here, keywords from the selected search results are appended to the query. A frequently used implementation of this approach is based on the *inverse document frequency* mentioned earlier in this section. Here, the underlying assumption is that more unique words are more valuable to the search. In this research we will use this strategy for selecting keywords to expand our queries with.

2.4 Vector-Based Feedback Strategies

Assuming vector representations of the query and search results are available, user feedback can also be applied to the queried vector directly. This is referred to as vector-based pseudo-relevance feedback. There are two commonly used methods for this [13]. First, the queried vector can be averaged with the positive search results. Second, the queried vector can be summed with the selected search results. Both strategies are implemented in our research.

A commonly used variation of averaging query vectors is based on Rocchio’s method for relevance feedback [13]. The high-level idea of this method is to move the query vector towards the selected vectors through assigning different weights to selected and queried vectors [20]. The version of Rocchio’s method implemented by most researchers today [3, 4] differs slightly from the original method, since it omits the negative feedback (i.e. non-selected documents) from the formula. As a result, this version of Rocchio can be seen as a weighted average between the (original) queried embedding and the (averaged) selected embeddings.

The weight of the queried embedding (α) and the weight of the averaged selected embeddings (β) can be set by a user. Still, the default/consensus values most research adheres to is $\alpha = 0.5$ and $\beta = 0.5$ [3]. Hence, we use Rocchio with those parameter values as a baseline method in our experiment.

3 Methodology

In this research we evaluate different relevance feedback strategies and text similarity methods with the objective of reducing review effort. As shown in Fig. 2, the method for accomplishing this is based on iteratively presenting the user with a set of (10) results to accept or decline. Thereafter, the accepted documents are used to improve the query (and consequently results) for the next iteration.




Fig. 2. Flowchart of the method.

For returning the results (i.e. relevance ranking) we experiment with different text similarity methods, levels of textual granularity, and ranking methods. These methods are explained in the first and second subsections. Next, for processing the feedback given after each iteration, we experiment with different feedback strategies. These are explained in the third and fourth subsections.

3.1 TF-IDF-Based Text Similarity

Our baseline method is based on TF-IDF, which vectorizes text based on multiplying the frequency of a word in a document with the *inverse* frequency of

that word in the entire data set. Hence, TF-IDF looks at the ‘uniqueness’ of a word instead of just the frequency. The result of TF-IDF is a sparse vector where all the words that exist in a document are given their TF-IDF value (words that do not appear in a specific document have zero values).

In our research we use TF-IDF to find similar documents using ‘More-Like-This’ (MLT). In summary, MLT queries the terms from a document (that have the highest TF-IDF values) individually using the document index. The documents returned by these queries are ranked based on combining their MLT scores, which is defined as the sum of the TF-IDF values of the matching terms between the queried document and the returned document.

Note, we are aware that a commonly used implementation of TF-IDF to identify textually similar texts is based on computing the cosine distance between the TF-IDF vectors. However, computing the cosine distance between n vectors results in quadratic time and space complexity ($O(n^2)$). Hence, we use MLT instead.

3.2 BERT-Based Text Similarity

Our semantic similarity method is based on Sentence-BERT (SBERT) embeddings. These embeddings can be used to find similar texts using ‘dense-vector search’ (DVS). The distance metric used for this is cosine similarity, which computes the similarity between two vectors (i.e. embeddings) based on the cosine value of the angle between the two vectors [18]. This angle is computed through dividing the dot-product (which is the sum of products from both vectors) by the length of the vectors. This means that the potential values of this similarity measure range from -1 (completely opposite) to 1 (completely similar). The formula for this is given below:

$$\text{cosine similarity} = \cos(\theta) = \frac{\mathbf{A}\mathbf{B}}{\|\mathbf{A}\|\|\mathbf{B}\|} = \frac{\sum_{i=1}^n \mathbf{A}_i\mathbf{B}_i}{\sqrt{\sum_{i=1}^n (\mathbf{A}_i)^2}\sqrt{\sum_{i=1}^n (\mathbf{B}_i)^2}} \quad (3)$$

In order to find embeddings with a high cosine similarity in a large dataset efficiently we use Hierarchical Navigable Small Worlds (HNSW) based vector search [15]. HNSW is an algorithm that finds the k most similar documents to a query with logarithmic time and space complexity ($O(\log(N))$).

3.3 Paragraph-Based Document Rankings

Because relevant information can be exclusive to a specific part of a document [1], we conduct experiments on two levels of text granularity: the document and paragraph levels. Here, the paragraph level means that we query and retrieve paragraphs instead of documents. However, for the relevance ranking we only consider documents. As a result, the experiments on the paragraph level require a document ranking to be derived from the returned paragraphs. For this, we define two different paragraph-based document rankings.

The first ranking method is based on taking the highest ranked paragraph of a document in the ranking as the overall document ranking. For example, say we return 6 paragraphs from 3 unique documents in the following order: $\{d_1, d_2, d_2, d_3, d_3, d_3\}$ where d_i refers to a paragraph from document i . Then, our document ranking will be as follows: $\{1, 2, 3\}$. Note how the ranking from the first paragraph of a document determines its position in the document ranking.

The second ranking method is based on counting the number of paragraphs per document in the ranking, and using that to rank the documents. For example, say we return the same 6 paragraphs in the following order: $\{d_1, d_2, d_2, d_3, d_3, d_3\}$ where d_i refers to a paragraph from document i . Then, our document ranking will be as follows: $\{3, 2, 1\}$. Note how in contrast to the previous ranking method, the number of paragraphs per document determines the ranking.

3.4 Baseline Feedback Strategies

We use three methods for the relevance feedback experiments. The first baseline method is to re-use the original queried embedding. If a user accepts/declines documents, the ranking of documents remains constant. This is referred to as ‘no feedback’. The second baseline method is based on keyword expansion. In keyword expansion, each iteration a number of keywords from the documents with *positive* pseudo-relevance feedback is selected and appended to the original query using an OR operator. As a result, the selected keywords serve as a pre-filter for the original query where the documents should contain at least 1 of the collected keywords to be considered. In our research, this selection is done based on the (highest) *inverse document frequency* value.

3.5 Vector-Based Feedback Strategies

Since the queried and collected texts are represented by vectors (BERT or TF-IDF), the relevance feedback methods are based on vector operations. These vector operations are based on summing and averaging the vectors. For this, we experiment with both cumulative (i.e. include the queried vector in the average/sum) and non-cumulative feedback (i.e. exclude the queried vector in the average/sum).

Also, for text similarity methods implemented on the paragraph level, relevance feedback can be amplified to the document level. If a given paragraph receives positive relevance feedback, then that feedback can be extended to other paragraphs that have the same parent document. In our research this will be referred to as ‘amplified feedback’ (or ‘amp’ in tabular formats). Note, feedback amplification is not applicable to any of our baseline methods.

4 Experimental Setup

This section provides an overview of the experiments and our data (and preprocessing). The first experiment focuses on comparing the performances of different text similarity methods. The second experiment focuses on implementing the best performing similarity method using different relevance feedback strategies.

4.1 Data and Preprocessing

Our research uses the RCV-1 v2 dataset in all experiments [2]. This dataset was made public in 2005 by Reuters News and consists of 806784 news articles. Due to hardware constraints, we did not use the complete dataset in our experiment. Instead, we randomly sampled 300 articles per topic. For this, we sampled 15 (unrelated) topics that are equally split in train, validation, and test. We call this data set the ‘default set’. Moreover, besides these topics we also sampled 4 topics that share the same parent topics. This set will be referred to as the ‘ambiguous’ set.

As for preprocessing, in TF-IDF we filter out stop words, numbers, and convert the text to lowercase. The stop word list used for this is publicly available on GitHub¹. For the SBERT-based experiments, we only remove numbers and special characters.

Finally, for the experiments on the paragraph level, we first split the text into sentences using the `<p>...</p>` tags in the XML files from RCV-1 v2 dataset [2]. Next, a paragraph is created through concatenating every 3 adjacent sentences of a document (and the remainder). For SBERT the number of words in a paragraph cannot exceed 384. Hence, we verified that the paragraphs in the topic sets do not exceed that limit.

4.2 Text Similarity Methods

For the text similarity methods, we compare two approaches: MLT (which is based on TF-IDF) and DVS (which is based on SBERT). For both methods, we iterate through our set using ‘query by document’ (QBD). Each time we query a document/paragraph to return other documents/paragraphs that belong to the same topic.

For MLT, we set two parameters. First, the minimum document frequency for words to be considered (minDf). Second, the maximum document frequency for words to be considered (maxDf). Due to limited computing resources, we didn’t conduct a full grid search to find the optimal values for these parameters. Instead, we conducted a manual search on the train and validation sets and used the average of the most performant parameter values on our default (and ambiguous) test set. Note, we conduct these experiments on the paragraph *and* document level.

For DVS, we do not have parameters to set. Hence these experiments are conducted directly on our default and ambiguous set. However, we do experiment with three different pre-trained SBERT models. These are selected based on being the general-purpose models in the SBERT documentation [19]. An overview of these models can be found in Table 1. Note, due to SBERT’s maximum context length, we conduct these experiments on the paragraph level only.

Finally, we evaluate the similarity methods based on (the macro averaged) recall, precision and F1 scores. For these metrics, the definition of a ‘true positive’

¹ <https://github.com/stopwords-iso/stopwords-en>.

Table 1. Selected BERT models’ characteristics according to [19].

Model name	Size	#Dimensions	Speed (sentences/sec)
all-mpnet-base-v2	420 MB	768	2800
all-MiniLM-L12-v2	120 MB	384	7500
all-MiniLM-L6-v2	80 MB	384	14200

Algorithm 1: Pseudo relevance feedback experiment

```

input : paragraph, maxRecall
output: Iterations needed to achieve recall
1 iterations = 0
2 while recall(acceptedDocuments)  $\leq$  maxRecall do
3   filter = acceptedDocuments + declinedDocuments
4   results = query(paragraph, filter)           // Returns top 10 results.
5   for result in results do
6     if feedback for result is positive then
7       paragraph = processFeedback(result)
8       acceptedDocuments += result
9     else
10      declinedDocuments += result
11  iterations += 1
12 return iterations

```

is a returned document that is of the same set as the queried document. This positive set is based on the annotations of the dataset. For example, if we query a document annotated as ‘sports’, then the document returned (as similar) should also be annotated as ‘sports’ to be considered a true positive.

4.3 Relevance Feedback

For our relevance feedback experiments we implement a form of pseudo-relevance feedback. The pseudo-code of this experiment can be found in Algorithm 1. In summary, the algorithm iteratively collects relevant documents until a pre-defined recall threshold is reached. For each iteration, the 10 most relevant documents are collected based on our selected text similarity method. Next, the relevant documents in this set are used to improve the query for the next iteration using relevance feedback.

Note, documents are considered relevant based on the same definition of a true positive as mentioned earlier. As a result, this experiment is based on pseudo-relevance feedback. Finally, for performance evaluation, we record the average number of iterations needed to achieve exactly 80% recall (which is a commonly used threshold in information retrieval [21]).

4.4 Configuration

The experiments are conducted on a local device. The CPU of this device is an Intel(R) Core(TM) i7-10610U CPU @ 1.80 GHz and it has 16 GB of RAM memory. As for software, the TF-IDF and DVS based text similarity experiments including MLT were conducted using Solr², which is built on top of Lucene³. Finally, the source code used in this experiment is publicly available on GitHub⁴.

5 Results

This section provides an overview of the results from our experiments. In the first subsection, we discuss the results of the individual text similarity methods. In the second subsection, we present the results of the different relevance feedback strategies.

5.1 Text Similarity Methods

For our TF-IDF-based approach, the manual search on the train and validation sets resulted in the parameter values of $\text{maxDf} = 0.8$ and $\text{minDf} = 0$ on both the paragraph and document level. Hence, the TF-IDF related results are based on these parameter values. Next, for the DVS experiments, *all-mpnet-base v2* was the best performing pre-trained SBERT model. Hence, the results of this model are shown in this section.




Fig. 3. Different configurations for the TF-IDF-based approach on the default set.

² <https://solr.apache.org/>.

³ <http://lucene.apache.org/java/340/scoring.html>.

⁴ <https://github.com/TimoKats/MasterThesis>.




Fig. 4. Different configurations for DVS on the default set.

For our default test set, the recall-precision graphs are shown in Fig. 3 and Fig. 4. In both figures, it is apparent that deriving a document ranking from its *highest* ranked paragraph outperforms ranking documents based on *counting* the paragraphs. Moreover, for both methods querying the first paragraph of a document slightly outperforms querying random paragraphs. Next, for the TF-IDF-based method specifically, the experiments on the document level outperform the experiments on the paragraph level.




Fig. 5. Identified configurations of DVS and TF-IDF-based approach on the ambiguous set.

Finally, when comparing the performance of the optimal configuration of both approaches, DVS (with a first-based ranking) outperforms all TF-IDF-based approaches. Next, when comparing the optimal configurations of the approaches on the *ambiguous* set, we again see that DVS outperforms the TF-IDF-based approach (see Fig. 5).




Fig. 6. Time consumed to create SBERT embeddings per paragraph.




Fig. 7. RAM% consumed to create SBERT embeddings per paragraph.

For all pre-trained SBERT models used in our experiments, RAM usage and computation time per embedding are shown in Fig. 6 and 7. In these figures, the x-axis refers to the individual paragraphs (i.e. each x-tick is a paragraph) and the y-axis refers to the time taken/memory consumed to create an embedding for that paragraph. Here, despite the large variance of results shown in Fig. 6, both these figures confirm the pattern described in the documentation from SBERT (shown in Table 1), which is that the larger models are more time and memory consumptive. Hence, the results of *all-mpnet-base v2* shown in this section do have extra computational costs.

5.2 Relevance Feedback

For the second set of experiments, we experimented with different feedback strategies using the identified text similarity method (DVS). For these experiments, the average number of iterations (and standard deviation) needed to achieve 80% recall for all methods and datasets is shown in Table 2. Note, every iteration translates to a review effort of 10 paragraphs. Similar to the previous experiments, the results are available on the default and ambiguous set.

First, the results on the default test set. Here, it is apparent that the feedback methods based on vector operations require fewer iterations and have a lower standard deviation than the baseline methods. The best performing feedback method is summing the vectors. Next, the results on the ambiguous set. Here, all similarity methods require more iterations to reach 80% recall and have a

Table 2. Iterations needed to achieve 80% recall (baselines are in the top cells, cumulative approaches are in the middle cells, non-cumulative approaches are in the bottom cells).

Feedback strategy	Mean (Default)	Std Dev (Default)	Mean (Ambiguous)	Std Dev (Ambiguous)
No feedback	33.36	8.84	46.77	10.04
Keyword expansion	32.24	8.03	44.23	9.94
Rocchio ($\alpha = 0.5, \beta = 0.5$)	29.50	3.79	37.98	5.11
Average	29.43	1.51	32.65	2.97
Average (amp)	30.17	1.03	32.99	1.80
Sum	28.29	0.75	29.41	2.13
Sum (amp)	28.46	0.50	30.54	1.31
Average	31.54	1.14	38.68	4.41
Average (amp)	32.81	1.08	38.41	2.64
Sum	32.80	1.12	38.80	5.28
Sum (amp)	32.20	0.87	38.20	3.77

higher standard deviation than they have on the default set. Still, feedback methods based on summing the vectors (cumulatively) gives the best results.

It is apparent that cumulative feedback methods outperform non-cumulative feedback methods for all vector-based relevance strategies. Moreover, the difference between averaging and summing vectors is smaller when using non-cumulative strategies.

A noteworthy finding when comparing the results from the default and ambiguous set is the gap between the minimal baseline method (of no feedback) and the optimal feedback method (of cumulatively summing the vectors). On the default set, the reduction of review effort (measured in the number of iterations to achieve 80% recall) is 17.85%. On the ambiguous set however, this reduction of review effort equals 59.04%. Another noteworthy finding is that amplifying the feedback to sibling paragraphs reduces the standard deviation, but not the number of iterations needed to achieve 80% recall.

Finally, for all methods we measured average time taken per iteration. The results for this are shown in Table 3. These results show that relevance feedback strategies based on vector operations (average and sum) add little latency to the experiment compared to no feedback. Still, amplifying feedback to sibling paragraphs does add some latency to the experiment for both averaging and summing vectors. However, keyword expansion adds the most latency to the experiments.

6 Discussion

This section is a discussion of the results. First, we interpret the results for our two main experiments. Second, we discuss the limitations of our research and the resulting suggestions for future work.

Table 3. Average execution times for different relevant feedback strategies.

Strategy	Average time per iteration (in seconds)
No feedback	0.02298
Average	0.02386
Sum	0.02417
Rocchio	0.02901
Sum (amp)	0.04982
Average (amp)	0.05097
Keyword expansion	0.12093

6.1 Text Similarity Methods

For all of text similarity methods, some common denominators emerge from the results. First, querying the first paragraph gives better results than querying a random paragraph. This coincides with findings from related work, since prior research found that the effectiveness of ‘query by document’ approaches depends on the prevalence of relevant terms in the queried text [27]. Combining this finding with our results, it is probable that the first paragraphs in the RCV-1 v2 dataset outperform random paragraphs because they have a higher prevalence of relevant terms. This property could be exclusive to the news articles used in our research, and therefore might not apply to the data used in other domains.

Another commonality between the methods is related to the paragraph-based document rankings. Here, we see that for all methods ranking documents based on the first paragraph in the ranking gives better results than querying documents based on counting their paragraphs in the ranking. Potentially, this could be related to differences in the number of paragraphs per document. If a document only has one paragraph, then it is always bound to be at the bottom of a count-based ranking, regardless of how related/similar that document is.

As for comparing the different text similarity methods, both sets of experiments show that DVS outperforms the TF-IDF-based approach. Considering our DVS approach is based on BERT instead of TF-IDF, this finding makes sense. Because, BERT’s bidirectional self-attention mechanism [8] enables it to understand more ambiguous and context-sensitive language properties. For example, the usage of synonymy, homonymy and referrals.

However, it must also be noted that our findings are based on a pre-trained SBERT model (*all-mpnet-base v2*) that had higher computational costs with regards to RAM usage (see Fig. 7) and computation time (see Fig. 6) than the other pre-trained models in our experiment.

6.2 Relevance Feedback

For the second experiment, the best performing text similarity method (DVS, with a first-based document ranking) was implemented for relevance feedback.

For both the default and the ambiguous set, the experiments show that relevance feedback methods based on cumulatively summing the vectors reduce review effort the most. Interestingly, this improvement does not seem to come at a computational cost. Since Table 3 shows that the execution time of these methods (without feedback amplification to sibling paragraphs) is fairly similar to our minimal baseline method (of no feedback). Note, when optimizing for a low standard deviation, amplifying feedback to sibling paragraphs is beneficial, which does add latency to the experiment.

6.3 Reflection on Implementation and Contributions

With regards to implementation, our results show that review effort can be decreased using text similarity-based relevance feedback methods. An important side note in this finding is that relevance feedback accomplishes this through only re-ranking documents. As a result, this strategy does not produce false negatives automatically without the awareness of the user (in contrast to an active learning based approach). This is particularly important in use-cases that require a high recall.

In terms of scientific novelty and contributions, we should state that the concept of relevance feedback has been studied before [28]. However, certain parts within our implementation are (to our knowledge) novel and therefore contribute to science. First, the evaluation of different paragraph-based document rankings contribute to the domain of paragraph-based document-to-document retrieval. Given the rise of large language models (that are generally limited to the paragraph level [14]) and the fact that relevant information can be exclusive to a specific part of a document [1], these findings are applicable beyond the technologies/embeddings used in this research. Second, our results show that the usage of sibling paragraphs in relevance feedback can reduce the standard deviation of review effort. To our knowledge, this technique and finding is novel. Moreover, a more ‘stable’ reduction of review effort could be favorable in real-world scenarios.

Hence, our findings are not just novel, but also applicable and practically implementable.

6.4 Limitations

First, data in real-world information retrieval scenarios tends to be heterogeneous (e.g. emails, documents, direct messages, etc.) [24]. But, the data used in this research is fairly homogeneous, since it only consists of news articles. Hence, there’s a slight mismatch between the type of data used in this experiment and the real-world data this technique could be applied to. Thus, the lack of data heterogeneity in this experiment is a limitation of this research.

Second, the size of the dataset (only 300 articles per topic) is smaller than most real-world information retrieval scenarios. Moreover, our data only consists of news articles. Therefore, certain findings in our research (e.g., the fact that querying the first paragraph slightly outperforms querying random paragraphs)

might not apply on other datasets. Still, given the fact that the RCV-1 v2 dataset is commonly used as a benchmark dataset in information retrieval [7], the results are still an adequate indication of our method’s performance.

6.5 Future Work

Our first suggestion for future work is related to the size of our dataset. As mentioned, due to computational constraints we only sampled 300 articles per topic. However, datasets in real-world information retrieval applications can be much larger than that. As a result, future work could run these methods on larger datasets to research their scalability.

Our second suggestion for future work is related to the level of data heterogeneity in the experiments. As mentioned in the previous subsection, data used in real-world information retrieval scenarios tend to be heterogeneous whereas the data used in this experiment is fairly homogeneous. As a result, future work could research to what extent heterogeneous data impacts the ability of different text similarity methods to return similar documents.

Next, this research uses Solr’s ‘MoreLikeThis’ functionality to increase the speed of our TF-IDF-based similarity experiments. Meaning, we didn’t use any vector space scoring to compute the similarities between the TF-IDF vectors. The reasoning behind this is that computing the (cosine) similarities between n TF-IDF vectors results in quadratic ($O(n^2)$) time and space complexity, which is simply not viable in a real-world application. However, recent innovations have made it possible to compute the pairwise similarities between (sparse) vectors much faster. An example of this is the ChunkDot Python library⁵, which splits the TF-IDF matrix into chunks and computes the similarities in parallel. Future work could use this innovation to experiment with TF-IDF-based cosine similarity as an additional text similarity method.

Finally, it will be interesting to further experiment with our relevance feedback strategies in combination with using large language models for the embeddings, or potentially also for pseudo or proxy feedback, both in the context of discovery and information retrieval tasks, as well as other tasks such as retrieval augmented generation.

7 Conclusion

This research evaluates the impact of changing the (text similarity-based) relevance rankings based on relevance feedback. For this, the first research question was formulated as follows: what text similarity methods are most suitable for relevance feedback? Our results show that the most suitable text similarity method for this is DVS (using BERT-based dense-vector representations) where the highest ranked paragraph determines the document ranking.

Next, the second research question was formulated as follows: to what extent can relevance feedback help to reduce review effort? We have outlined a variety

⁵ <https://pypi.org/project/chunkdot/>.

of relevance feedback strategies. Our results show that (compared to processing no relevance feedback) our best-scoring the relevance feedback strategy reduces review effort between 17.85% and 59.04%, given a target recall level of 80%.

Given the recall-oriented nature of many information retrieval applications [23], the results for the relevance feedback experiments are encouraging. Since, in contrast to an active learning based strategy (which typically used for this purpose), this approach reduces review effort through only re-ranking documents. As a result, there are no false negatives created without the awareness of the user.

References

1. Adebayo, K.J.: Multimodal legal information retrieval. Ph.D. thesis, University of Luxembourg, Luxembourg (2018)
2. Amini, M.-R., Goutte, C.: Reuters RCV1 RCV2 multilingual, multiview text categorization test collection. UCI Machine Learning Repository (2013)
3. Arampatzis, A., Peikos, G., Symeonidis, S.: Pseudo relevance feedback optimization. *Inf. Retrieval J.* 269–297 (2021)
4. Cai, T., et al.: Scalable relevance ranking algorithm via semantic similarity assessment improves efficiency of medical chart review. *J. Biomed. Inform.* 104109 (2022)
5. Chai, C.P.: Comparison of text preprocessing methods. *Nat. Lang. Eng.* 509–553 (2023)
6. Chandrasekaran, D., Mago, V.: Evolution of semantic similarity—a survey. *ACM Comput. Surv.* 1–37 (2022)
7. Deolalikar, V.: How valuable is your data? A quantitative approach using data mining. In: 2015 IEEE International Conference on Big Data (Big Data), pp. 1248–1253. IEEE (2015)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, vol. 1, pp. 4171–4186. Association for Computational Linguistics, Minneapolis (2019)
9. Gorman, J., Curran, J.R.: Scaling distributional similarity to large corpora. In: *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pp. 361–368 (2006)
10. Grossman, M.R., Cormack, G.: Continuous active learning for TAR. *Journal* 1–7 (2016)
11. Jaccard, P.: Nouvelles recherches sur la distribution florale. *Bull. Soc. Vaud. Sci. Nat.* 223–270 (1908)
12. Joshi, S., Contractor, D., Ng, K., Deshpande, P.M., Hampp, T.: Auto-grouping emails for faster e-discovery. *Proc. VLDB Endow.* 1284–1294 (2020)
13. Li, H., Mourad, A., Zhuang, S., Koopman, B., Zuccon, G.: Pseudo relevance feedback with deep language models and dense retrievers: successes and pitfalls. *ACM Trans. Inf. Syst.* 1–40 (2023)
14. Liu, Y., et al.: Summary of ChatGPT-related research and perspective towards the future of large language models. *Meta-Radiol.* 1(2), 100017 (2023)

15. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs, pp. 824–836. IEEE Computer Society (2020)
16. Manning, C.D.: Introduction to Information Retrieval. Syngress Publishing (2008)
17. Piramuthu, O.B.: Multiple choice online algorithms for technology-assisted reviews. In: Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing, pp. 639–645 (2023)
18. Rahutomo, F., Kitasuka, T., Aritsugi, M.: Semantic cosine similarity. In: The 7th International Student Conference on Advanced Science and Technology ICAST, p. 1 (2012)
19. Reimers, N., Gurevych, I.: Sentence-BERT: sentence embeddings using Siamese BERT-networks. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) 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. 3982–3992. Association for Computational Linguistics, Hong Kong (2019)
20. Rocchio, J.J.: Document retrieval system-optimization and evaluation. In: DIR 2009 Dutch-Belgian Information Retrieval Workshop, p. 99 (2009)
21. Roitblat, H.L.: Probably reasonable search in ediscovery (2022). [arXiv:2201.12376](https://arxiv.org/abs/2201.12376)
22. Sánchez, D., Batet, M.: A semantic similarity method based on information content exploiting multiple ontologies. *Expert Syst. Appl.* 1393–1399 (2013)
23. Song, J.J., Lee, W., Afshar, J.: An effective high recall retrieval method. *Data Knowl. Eng.* **123**, 101603 (2019)
24. Sperling, M., Jin, R., Rayvych, I., Li, J., Yi, J.: Similar document detection and electronic discovery: so many documents, so little time. In: ICAIL 2013 Workshop on Standards for Using Predictive Coding, Machine Learning, and Other Advanced Search and Review Methods in E-Discovery (DESI V Workshop), Rome, Italy (2013)
25. Teernstra, L., van der Putten, P., Noordegraaf-Eelens, L., Verbeek, F.: The morality machine: tracking moral values in tweets. In: Boström, H., Knobbe, A., Soares, C., Papapetrou, P. (eds.) IDA 2016. LNCS, vol. 9897, pp. 26–37. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46349-0_3
26. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems (2017)
27. Yang, E., Lewis, D.D., Frieder, O., Grossman, D.A., Yurchak, R.: Retrieval and richness when querying by document. In: Biennial Conference on Design of Experimental Search & Information Retrieval Systems, pp. 68–75 (2018)
28. Zhang, H., Cormack, G.V., Grossman, M.R., Smucker, M.D.: Evaluating sentence-level relevance feedback for high-recall information retrieval. *Inf. Retrieval J.* 1–26 (2020)