

UNIMIB at TREC 2021 Clinical Trials Track

Georgios Peikos, Oscar Espitia, and Gabriella Pasi

Information and Knowledge Representation, Retrieval, and Reasoning (IKR3) Lab,
Department of Informatics, Systems, and Communication (DISCo),
University of Milano-Bicocca, Milan, Italy
{georgios.peikos, oscar.espitiamendoza, gabriella.pasi}@unimib.it

Abstract. This contribution summarizes the participation of the UNIMIB team to the TREC 2021 Clinical Trials Track. We have investigated the effect of different query representations combined with several retrieval models on the retrieval performance. First, we have implemented a neural re-ranking approach to study the effectiveness of dense text representations. Additionally, we have investigated the effectiveness of a novel decision-theoretic model for relevance estimation. Finally, both of the above relevance models have been compared with standard retrieval approaches. In particular, we combined a keyword extraction method with a standard retrieval process based on the BM25 model and a decision-theoretic relevance model that exploits the characteristics of this particular search task. The obtained results show that the proposed keyword extraction method improves 84% of the queries over the TREC’s median NDCG@10 measure when combined with either traditional or decision-theoretic relevance models. Moreover, regarding RPEC@10, the employed decision-theoretic model improves 85% of the queries over the reported TREC’s median value.

1 Introduction

The TREC 2021 Clinical Trials Track¹ focuses on the task of retrieving clinical trials that are eligible for a particular patient. Specifically, patient-related information is provided as a query consisting of a synthetic case in the form of an admission note. Each query describes a patient by means of certain conditions and observations, and, in general, queries are longer than those in traditional ad-hoc retrieval tasks. Moreover, clinical trials, i.e., the considered document collection, are multi-fielded documents, where a core aspect is constituted by the fields that specify the inclusion and exclusion criteria. Those are not all-inclusive statements about the trial to the point that other trial information can be ignored, but they are essential aspects for defining a trial’s eligibility (relevance) to a given patient.

This report outlines our team’s (UNIMIB) submission to TREC 2021 Clinical Trials Track, it discusses the experimental setup, and compares our results to the reported TREC’s median performance. In particular, we examined the effect of different query representations for retrieving eligible clinical trials. Moreover, we combined those different query representations with several relevance models, one of which is a novel decision-theoretic model that leverages the peculiarities of this search task in the relevance estimation.

2 Methodology

This section describes the three components of the submitted runs, i.e., the query and document representations, along with the employed relevance models. The first one is the BM25 model that leverages an ad-hoc query representation. In addition, the second relevance model is a decision-theoretic model whose properties are based on the peculiarities of this particular search task; finally, a neural network-based model in a re-ranking setup has been employed.

¹ <http://www.trec-cds.org/2021.html>

2.1 Query representation

The provided queries consist of synthetic patient cases created by medical experts in the form of admission notes. Previous studies [1, 2] have investigated the effect of different query representations on the system’s performance. Their findings suggested that using the whole verbose query is not as effective as using its ad-hoc representation. Therefore, given a query that is a verbose representation of a synthetic patient case, we extracted keywords that better represent its content.

To this aim, we have employed EmbedRank++ proposed in [3], which is a fast, unsupervised, and collection-independent method for extracting key phrases and keywords; these advantages make this method suitable for extracting keywords from verbose queries. Firstly, given a query, we estimated its embedding representation, and then we estimated the embedding representation of every of its tokens, both unigrams and bigrams. Having done that, we selected the keyword or key phrase that is most similar to the query by measuring their cosine similarity. Then, by employing the maximal marginal relevance (MMR) formula, as introduced in [3], we iteratively selected new keywords or key phrases similar to the query and different from the already extracted keywords. The employed keyword extraction approach requires the setting of two parameters; the total number of extracted keywords, and a parameter λ that balances the trade-off between similarity to the query and diversity between the extracted keywords. Concluding, by using this tool, we were able to extract those keywords that better represent each query, and at the same time, avoid redundant terms.

For the submitted runs, we have used two different query representations, i.e., the original verbose query Qd , provided by the organizers, and an ad-hoc representation that contains the extracted keywords, Qk .

2.2 Document Representation

In the context of this particular search task, as pointed out in Section1, several document sections have a different effect on the relevance of the whole document to the query. Specifically, without underestimating the contribution of the other document sections, the parts that refer to a trial’s inclusion and exclusion criteria in the eligibility section are of great importance.

Besides the eligibility section, other important document sections are its title, description, summary, and studied condition. Consequently, three of our submitted runs leverage an index created by combining the eligibility section with all the document sections mentioned above; throughout this paper, this index will refer to as L_{comb} . For the other two runs the index was created for only the title and summary with inclusion criteria when it is specified, and successfully extracted, this index will be referred to as L_{comb}^* .

Two additional indices have been created for each document, namely L_{in} which indexes only the trial’s inclusion criteria, and L_{ex} which indexes only the trial’s exclusion criteria. For their extraction, several regex rules that leverage their semi-structured nature have been used. In the cases where the extraction of the inclusion and exclusion criteria was not feasible, the whole eligibility section has been indexed for both indices, i.e., the L_{in} and L_{ex} . In addition to those two indices, a third one has been created. Here, the title, description, studied condition, and the summary sections were combined in a single text and indexed; this index will be referred to as L_{main} .

In summary, the first, fourth and fifth submitted run uses the combined document representations L_{comb} and L_{comb}^* to estimate the relevance of a clinical trial, while the relevance decision for the second and third submissions rely on the L_{in} , the L_{ex} and the L_{main} indices, and will be further analyzed in section 2.5

2.3 Retrieval Models

Given the various query and document representations, we have combined them with several relevance models. The considered models are the traditional BM25 model, a neural network-based model in a re-ranking setup, trained on previous clinical trials collections, and a decision-theoretic relevance model, which we have introduced for this particular search task. Specifically, the decision-theoretic relevance model is based on the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS) [4].

2.4 Standard Approach

A traditional retrieval approach has been employed that leverages the Qk representation along with the Lcomb index and the BM25 model to estimate an overall relevance value and rank the documents.

2.5 Decision-Theoretic Approach

As stated previously, the trial’s inclusion and exclusion criteria have a significant role in the relevance (eligibility) of a trial to a particular patient. Specifically, we argue that the semantic similarity of a patient health record, in our case the query, to a trial’s inclusion criteria contributes positively to the trial’s overall relevance. In contrast, the semantic similarity to its exclusion criteria contributes negatively. In addition, we argue that the semantic similarity of a query to the remaining trial’s sections, i.e., its title, description, summary, studied condition, also contributes positively to its overall relevance. These are task-specific characteristics that standard retrieval approaches usually do not consider, while a decision-theoretic model allows us to incorporate them in the relevance estimation.

In detail, to estimate the overall trial’s relevance based on the previous assumptions, firstly, we have employed three different formal representations, namely the *Lin*, the *Lex* and the *Lmain*, that have been introduced in Section 2.2. Then, given these three representations, we estimated, for each query, the topical relevance for each of them by using the BM25 model. By doing that, we ended up with three relevance scores that have a different contribution, either positive or negative, to the document’s overall relevance. The final step was aggregating those scores by explicitly considering their independent contribution to relevance. That has been achieved by incorporating TOPSIS, a multi-criteria decision analysis method that can provide a final document ranking based on the task-specific characteristics.

Specifically, the implementation of TOPSIS requires several components to be considered:

- A set of alternatives that have to be ranked, in our case the documents.
- A set of considered criteria, with their related performance scores (in our case the three relevance scores), *R_{in}*, *R_{ex}*, *R_{main}* calculated w.r.t. the different document representations using BM25. In order for TOPSIS to be applied, those scores must be monotonically increasing or decreasing and normalized.
- A weighting vector associated with the importance of each criterion to the final relevance decision.

In addition to the components mentioned above, TOPSIS allows to specifically state the contribution, positive or negative, of a criterion to the final overall score. In this particular search task, the *R_{in}* and *R_{main}* criteria contribute positively, therefore they are *benefit criteria*, while the *R_{ex}* contributes negatively, and, as a result, this criterion is a *cost criterion*.

TOPSIS is based on the concept that the chosen alternative, i.e., a document, should have the shortest geometric distance from a positive ideal solution and the longest from a negative ideal solution. In general, the positive ideal solution is a vector consisting of the maximum values of the benefit criteria and the minimum values of the cost criteria. In contrast, the negative ideal solution consists of the maximum values of the cost criteria and the minimum values of the benefit criteria.

Therefore, given the decision matrix related to a specific query containing the normalized performance (relevance) scores referred to the considered sections of each document, the weighted average of the matrix can be computed over the three considered criteria, to obtain the final relevance score. Moreover, the positive and negative ideal solutions can be generated as follows: the positive ideal solution is represented by a vector that contains the highest scores of the *R_{in}* and *R_{main}* criteria and the lowest score of the *R_{ex}* criterion, present in the normalized decision matrix, i.e., across all the documents for a given query; the negative ideal solution is the exact opposite.

Finally, for each query-document pair, the euclidean distance of the document from the positive and negative ideal solutions is measured. The documents are then ranked based on the relative closeness of their overall relevance score to the ideal solution.

2.6 Neural Approach

For two of the runs, namely UNM_4 and UNM_5, we implemented a re-ranking approach based on BM25 and Bert. BM25 employs the query representation Qd and the document representation $Lcomb^*$; on the re-ranking side, Bert for sentences classification was fine-tuned in a low data regime by using collections from previous TREC Biomedical tracks ²; and also the collection indicated by the organizers [2]. Concerning the document representation, we also used $Lcomb^*$. In this way, the combination of query and document does not exceed the limit of tokens of the transformer network. The training procedure was adopted by following previous approaches [5].

3 Experimental Setup

Several steps of our methodology required some parameter selection. Therefore, this section presents a few details related to implementation aspects.

3.1 Keyword Extraction Method

To begin with, the keyword extraction method applied to the queries has required the application of a specific model for obtaining both document and word embeddings. Due to its simplicity and speed, we have used the DistilBERT proposed in [6]. Moreover, the length of the provided queries, in terms of tokens, was smaller than the model’s maximum input, and as a result, no action was needed.

In addition, we needed to set two more parameters related to the diversity between the selected keywords and the size of the new query. Due to the assumption that each query represents a different patient and that some conditions may be more complicated than others (therefore reflected in longer descriptions) we have modified the original implementation of the EmbedRank++ method. In particular, instead of extracting a constant number of key phrases from each query, we estimate the size of the new query dynamically. The length of the new query is equal to half the size of the initial query; stop-words have been removed. Regarding the second parameter, which is related to the diversity between the candidate keywords, it was set equal to 0.5 to capture many different aspects of each query.

3.2 Indexing & Topical Relevance

To index the collection and estimate the topical relevance values with respect to each of different document representations (indices), we have employed PyTerrier [7] and its BM25 implementation, respectively. For preprocessing, we have used the standard PyTerrier’s pipeline, i.e., porter-stemming and stopwords removal.

3.3 TOPSIS

As stated in Section 2.5, TOPSIS associates a weight with each considered criterion based on which each document is evaluated. These weights can be either learned from a portion of training data, manually set to specific values, or estimated with some weighting method. For the submitted runs R2 and R3, the weights have been set to fixed values equal to 0.33, the mean weight value of the three criteria, thus assuming a same importance for the three criteria. We refer to this model as *TT_MW* model.

3.4 Submitted Runs

Table 1 presents the submitted runs along with the corresponding components that have been analyzed in the previous sections.

² <http://www.trec-cds.org/>

Run ID	Run Name	Query Representation	Document Index	Relevance Model
R1	IKR3_BSL	Qk	L.comb	BM25
R2	IKR3_TT_MW_d	Qd	L.in, L.lex, L.main	TT_MW
R3	IKR3_TT_MW_k	Qk	L.in, L.lex, L.main	TT_MW
R4	UNM_4	Qd	L.comb*	BM25+BERT
R5	UNM_5	Qd	L.comb*	BM25+BERT

Table 1. A summary of the components employed in each submitted runs.

4 Results

Table 2 presents the results obtained from the five submitted runs, while Table 3 reports the number of queries that were improved compared to their reported TREC’s median performance. Thus, it can be seen that the mean effectiveness of our runs is well-above the averaged TREC’s median for three different measures, and also that the submitted runs improve the majority of the queries.

Measure	TREC’s Median	R1	R2	R3	R4	R5
NDCG@10	.304	.478	.431	.449	.252	.346
PREC@10	.161	.265	.255	.281	.152	.173
Reciprocal Rank	.294	.486	.452	.459	.289	.373
NDCG@5	Not Reported	.504	.458	.468	.278	.380

Table 2. Overall comparison with the across query averaged median values, where R1-5 represent the effectiveness obtained by our runs.

Furthermore, one can observe that the highest PREC@10 value was obtained by the R3 run that involves the TOPSIS model and the Qk representation. Also, the R2 run that exploits the Qd query representation, that is usually more noisy, yields results relatively close to the runs that employ the Qk representation.

Runs	R1	R2	R3
NDCG@10	63	59	58
PREC@10	63	62	64
Reciprocal Rank	53	53	53

Table 3. Number of improved queries over the TREC’s Median.

Finally, the fact that the associated criteria weights were manually set, instead of learned using a training query set, may have underestimated the effectiveness of the employed decision-theoretic model.

5 Conclusions

In this paper we have presented the results of our participation in the TREC 2021 Clinical Trials Track. Our objective was to investigate the effectiveness of several query representation combined with different retrieval models in this particular search task. The analysis of the results confirmed the effectiveness of the query representation approach (run R1). However, we also proved that a decision-theoretic model (run R2), yields similar results compared to a standard retrieval approach that is associated with an ad-hoc query representation (run R1). In addition, the decision-theoretic

model was able to improve the system’s precision when combined with a better query representation (run R3). Regarding the neural-based approaches, their performance was comparable to the reported TREC’s median.

Concluding, it can be seen that using a decision-theoretic model is beneficial for this particular search task. The decision-theoretic model almost met the performance of a standard approach without leveraging the use of the query preprocessing method. At the same time, it outperformed the standard approach in terms of PREC@10 when combined with the employed query preprocessing method. However, the fact that the criteria weights were manually set may have underestimated the model’s effectiveness; therefore, further research will be conducted in this direction.

6 Acknowledgements

This work was supported by the EU Horizon 2020 ITN/ETN on Domain Specific Systems for Information Extraction and Retrieval – DoSSIER (H2020-EU.1.3.1., ID: 860721).

References

1. Bevan Koopman, Liam Cripwell, and Guido Zuccon. Generating clinical queries from patient narratives: A comparison between machines and humans. In Noriko Kando, Tetsuya Sakai, Hideo Joho, Hang Li, Arjen P. de Vries, and Ryen W. White, editors, *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7-11, 2017*, pages 853–856. ACM, 2017.
2. Bevan Koopman and Guido Zuccon. A test collection for matching patients to clinical trials. In Raffaele Perego, Fabrizio Sebastiani, Javed A. Aslam, Ian Ruthven, and Justin Zobel, editors, *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*, pages 669–672. ACM, 2016.
3. Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl, and Martin Jaggi. Simple unsupervised keyphrase extraction using sentence embeddings. In Anna Korhonen and Ivan Titov, editors, *Proceedings of the 22nd Conference on Computational Natural Language Learning, CoNLL 2018, Brussels, Belgium, October 31 - November 1, 2018*, pages 221–229. Association for Computational Linguistics, 2018.
4. Ching-Lai Hwang and Kwangsun Yoon. *Multiple Attribute Decision Making: Methods and Applications - A State-of-the-Art Survey*, volume 186 of *Lecture Notes in Economics and Mathematical Systems*. Springer, 1981.
5. Maciej Rybinski, Jerry Xu, and Sarvnaz Karimi. Clinical trial search: Using biomedical language understanding models for re-ranking. *Journal of Biomedical Informatics*, 109:103530, 2020.
6. Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108, 2019.
7. Craig Macdonald and Nicola Tonellotto. Declarative experimentation in information retrieval using pyterrier. In *Proceedings of ICTIR 2020*, 2020.