

# On the Interpolation of Contextualized Term-based Ranking with BM25 for Query-by-Example Retrieval

Amin Abolghasemi  
m.a.abolghasemi@liacs.leidenuniv.nl  
Leiden University  
Leiden, Netherlands

Arian Askari  
a.askari@liacs.leidenuniv.nl  
Leiden University  
Leiden, Netherlands

Suzan Verberne  
s.verberne@liacs.leidenuniv.nl  
Leiden University  
Leiden, Netherlands

## ABSTRACT

Term-based ranking with pre-trained transformer-based language models has recently gained attention as they bring the contextualization power of transformer models into the highly efficient term-based retrieval. In this work, we examine the generalizability of two of these deep contextualized term-based models in the context of query-by-example (QBE) retrieval in which a seed document acts as the query to find relevant documents. In this setting — where queries are much longer than common keyword queries — BERT inference at query time is problematic as it involves quadratic complexity. We investigate TILDE and TILDEv2, both of which leverage BERT tokenizer as their query encoder. With this approach, there is no need for BERT inference at query time, and also the query can be of any length. Our extensive evaluation on the four QBE tasks of SciDocs benchmark shows that in a query-by-example retrieval setting TILDE and TILDEv2 are still less effective than a cross-encoder BERT ranker. However, we observe that BM25 could show a competitive ranking quality compared to TILDE and TILDEv2 which is in contrast to the findings about the relative performance of these three models on retrieval for short queries reported in prior work. This result raises the question about the use of contextualized term-based ranking models being beneficial in QBE setting. We follow-up on our findings by studying the score interpolation between the relevance score from TILDE (TILDEv2) and BM25. We conclude that these two contextualized term-based ranking models capture different relevance signals than BM25 and combining the different term-based rankers results in statistically significant improvements in QBE retrieval. Our work sheds light on the challenges of retrieval settings different from the common evaluation benchmarks. It could be of value as future work to study other contextualized term-based ranking models in QBE settings.

## CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking; Evaluation of retrieval results.**

## KEYWORDS

Query-by-example retrieval, term-based retrieval, Transformer models, BERT-based ranking

## ACM Reference Format:

Amin Abolghasemi, Arian Askari, and Suzan Verberne. 2022. On the Interpolation of Contextualized Term-based Ranking with BM25 for Query-by-Example Retrieval. In *Proceedings of the 2022 ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR '22)*, July 11–12, 2022, Madrid, Spain. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539813.3545133>

## 1 INTRODUCTION

Query-by-Example (QBE) retrieval is an Information Retrieval (IR) setting in which a seed document<sup>1</sup> acts as the query to represent the user's information need and the retrieval engine searches over a collection of the same type of documents [1, 19, 20, 29]. This retrieval setup is typical in professional, domain-specific tasks such as legal case law retrieval [1, 2], patent prior art search [11, 24, 25], and scientific literature search [1, 19, 20]. While using a document as a query could become challenging due to its length and complex semantic structure, prior work has shown that traditional term-based retrieval models like BM25 [27] are highly effective when used in QBE retrieval [1, 2, 28].

Recently, deep contextualized term-based retrieval models have gained attention as they bring the contextualization power of the pre-trained transformer-based language models into the highly efficient term-based retrieval. Examples of such models are DeepImpact [18], SPLADE [10], SPLADEv2 [9], TILDE [34], TILDEv2 [33], COIL [12], and uniCOIL [16]. Here, we specifically investigate TILDE, which is a term independent likelihood model, and its follow-up TILDEv2 which is a deep contextualized lexical exact matching model.

TILDE and TILDEv2, which are introduced as term-based re-ranking models, follow a recent paradigm in term-based retrieval where term importance is pre-computed with scalar term weights. Besides, to predict the relevance score, both of these models use the BERT tokenizer as their query encoder which means that they do not need to perform any BERT inference at query time to encode the query. However, leveraging tokenizer-based encoding of the query trades off the query representation and therefore effectiveness with higher efficiency at inference time [33]. While the effectiveness of these models is evaluated on tasks and benchmarks where we have short queries, e.g., MSMARCO Passage Ranking [21] and the TREC DL Track [7], in this paper, we evaluate them in the aforementioned QBE retrieval setting where queries are much longer than common keyword queries. In this regard, we address the following research questions:

RQ1 How effective are TILDE and TILDEv2 in query-by-example retrieval?

<sup>1</sup>Throughout this paper, we use the term “document” to refer to a unit of retrieval [17]



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICTIR '22, July 11–12, 2022, Madrid, Spain.

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9412-3/22/07.

<https://doi.org/10.1145/3539813.3545133>

A specific direction in answering *RQ1* is to investigate the ranking quality of TILDE and TILDEv2 in comparison with the effective cross-encoder BERT ranker [1, 22], which is described in section 2.4. We are interested in this direction for two reasons. First, the cross-encoder BERT ranker exhibits quadratic complexity in both space and time with respect to the input length [17] and this is aggravated in QBE where we have long queries. TILDE and TILDEv2, however, do not need any BERT inference at query time. Second, due to the maximum input length of BERT, cross-encoder BERT ranker, which uses the concatenation of the query and the document, might not cover the whole query and document tokens in a QBE setting, whereas in TILDE and TILDEv2, the query can be of any length and documents are covered up to the maximum length of BERT.

Additionally, since TILDEv2 pre-computes the term weights only for those tokens existing in the documents, one risk is that it might aggravate the vocabulary mismatch problem. A typical approach to address this issue is to use document expansion methods. Zhuang and Zuccon [33] use TILDE as their document expansion model for TILDEv2. We adopt that approach for our task and further investigate the impact of token-based document expansion with TILDE on the ranking quality of TILDEv2 in a QBE retrieval setting.

Apart from comparing TILDE and TILDEv2 to the cross-encoder BERT ranker, we also make a comparison to traditional lexical matching models (BM25 and Probabilistic Language models), which have been shown as strong baselines on QBE tasks in prior work [2, 28]:

*RQ2* What is the effectiveness of traditional lexical matching models with varying tokenization strategies in comparison to TILDE and TILDEv2?

To answer *RQ2* we will investigate the effect of using the BERT tokenizer [8] as pre-processing for traditional term-based retrieval models. By doing so, we are aligning the index vocabulary of traditional models with that of TILDE and TILDEv2, which could make our comparison more fair.

We will see in the Section 4 that BM25 shows a competitive ranking quality in comparison to TILDE and TILDEv2 in our QBE benchmark. Because of the similar quality on average, we are interested to see if the relevance signals of TILDE and TILDEv2 are different from that of BM25, to find out if the methods are complementary to each other. To this aim, we will investigate the following research question:

*RQ3* To what extent do TILDE and TILDEv2 encode a different relevance signal from BM25?

To address the question above, as it is described in details in Section 3.3, we will analyze the effect of the interpolation of the scores of TILDE and TILDEv2 with BM25.

Since TILDE and TILDEv2 are introduced as re-ranking models, we use four different tasks from the SciDocs evaluation benchmark [5] as a domain-specific QBE benchmark. This benchmark uses scientific paper abstracts as the query and documents. The retrieval setting in these tasks suits as a re-ranking setup because of the number of documents to be ranked for each query. Since that we are working in a domain-specific evaluation setting, we will also address the following research question:

*RQ4* To what extent does a highly tailored domain-specific pre-trained BERT model affect the effectiveness of TILDE and TILDEv2 in comparison to a BERT<sub>base</sub> model?

In summary our main contributions in this work are three-fold:

- We show that two recent transformer-based lexical models (TILDE and TILDEv2) are less effective in Query-by-Example retrieval than was expected based on results reported for ad hoc retrieval. This indicates that QBE retrieval is structurally different from other IR settings and requires special attention for methods development;
- We show that the relevance signals of TILDE and TILDEv2 can be complementary to that of BM25 as interpolation of the methods leads to an improvement in ranking effectiveness;
- We also investigate interpolations of BM25 with TILDE and TILDEv2 in an ideal setting where the optimal interpolation weight is known a priori, and by doing so, we show that more stratified approaches for the interpolation could result in higher gains from the interpolation of BM25 with TILDE and TILDEv2.

In section 2 we describe the retrieval models used in this work. In section 3 we provide details about our methods and experiments and in section 4 we analyze the results and discuss the answers to our research questions. Section 5 is dedicated to further analysis of the results, and finally, in Section 6 we provide the conclusion.

The code used in this paper is available at: <https://github.com/aminenv/lexica>

## 2 BACKGROUND: RETRIEVAL MODELS

In this section, we briefly introduce the retrieval models that we implement and evaluate in our experiments.

### 2.1 Traditional lexical matching models

*BM25*. For BM25 [27], we use the implementation by Elasticsearch<sup>2</sup> with the parameters  $k = 2.75$ , and  $b = 1$ , which was tuned over the validation set.

*Probabilistic Language Models*. For language modeling (LM) based retrieval [4, 13, 26], we use the built-in similarity functions of Elasticsearch for the implementation of language model with Jelinek Mercer (JM) smoothing [32].

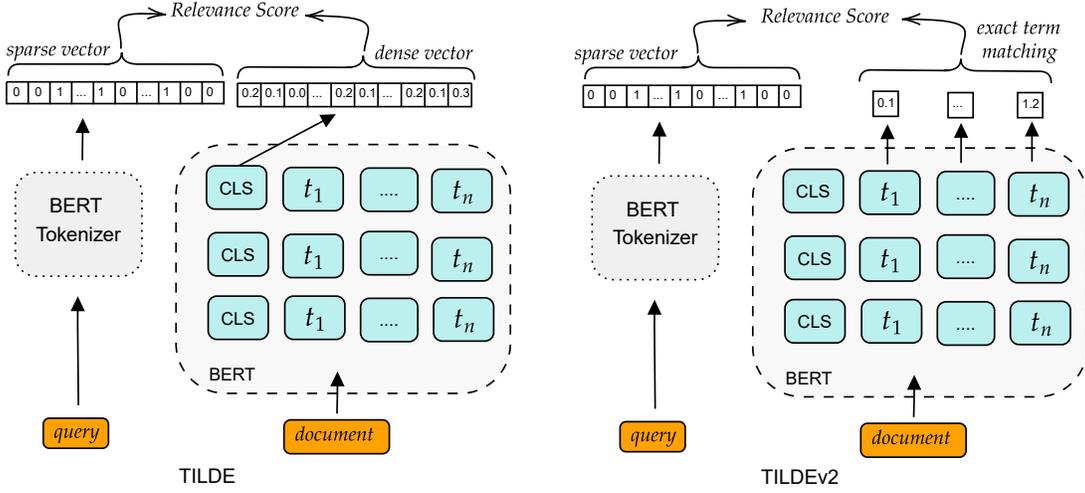
### 2.2 Term Independent Likelihood Model: TILDE

TILDE is a tokenizer-based term-based retrieval model which follows a term independence assumption and formulates the likelihood of a query as follows:

$$\text{TILDE-QL}(q|d) = \sum_i^{|q|} \log(P_\theta(q_i|d)) \quad (1)$$

in which  $q$  is the query, and  $d$  is the document. As Figure 1 shows, to compute the relevance score, the text of a document  $d$  is fed as the input for BERT and the log probability for each token is estimated by using a language modeling head on top of the BERT [CLS] token output. In other words, we are pre-computing the term weights over the complete BERT vocabulary. During both

<sup>2</sup><https://github.com/elastic/elasticsearch>



**Figure 1: Model architectures. Left: TILDE [34]. Right: TILDEv2 [33]. Both TILDE and TILDEv2 leverage the BERT tokenizer as their query encoder.  $t_i$  stands for the  $i$ th token of the document. The dense vector and sparse vector have the same length as the BERT vocabulary size.**

training and inference time, the query text is tokenized by using a BERT tokenizer and the resulting token IDs are used to look up the corresponding log probability from the likelihood distribution predicted in the output of the language modeling head. It is worth mentioning that the document likelihood can be computed in a similar way by swapping the query and document; however, we only use the query likelihood (Equation 1) in our experiments.

For TILDE, we use the implementation from the authors' code repository.<sup>3</sup> We report results for the TILDE model with different initial checkpoints as the BERT encoder for our fine-tuning procedure. TILDE<sub>BERT</sub> uses bert-base-uncased, TILDE<sub>SciBERT</sub> uses SciBERT, and TILDE<sub>MSMARCO</sub> uses a TILDE which is already fine-tuned on MSMARCO; we use TILDE<sub>MSMARCO</sub> in a zero-shot setting on our data.

### 2.3 Lexical Exact Matching: TILDEv2

TILDE has a drawback in which it expands each document to the size of BERT tokenizer vocabulary. To tackle this problem, the authors proposed TILDEv2. TILDEv2, which builds upon uniCOIL [16] and TILDE, follows a recent paradigm in contextualized lexical exact matching in which BERT is used to output a scalar importance weight for document tokens [16, 33]. As it is shown in Figure 1, in TILDEv2, the token representation is downsized into a scalar weight and the relevance score between a query and a document pair is computed by a sum over the contextualized term weights for all terms appearing in both query and document:

$$s(q, d) = \sum_{q_i \in q} \max_{q_i = d_j} (c(q_i) \times v_j^d) \quad (2)$$

Here,  $q$  and  $d$  are the query and the document respectively;  $d_j$  is the  $j$ th token of the document;  $v_j^d$  is the term importance weight for the  $j$ th token of  $d$ , and  $c(q_i)$  is the count of the  $i$ -th unique token which

is achieved by using the BERT tokenizer as the query encoder. In this equation,  $v_j^d$  is computed using the same method as in Lin and Ma [16] in which a *RELU* function is used on the projection layer to force the model to map the token representations into a positive scalar weight:

$$v_j^d = \text{ReLU}(W_{proj}^{1 \times n} \text{BERT}(d_j) + b) \quad (3)$$

in which  $d_j$  is the  $j$ th token in document  $d$  and  $b$  is the learnable bias parameter of the projection layer  $W_{proj}$ . Lin and Ma [16] show that using a scalar weight as term importance (uniCOIL [16]) instead of a vector representation (COIL [12]) results in a decrease in the effectiveness; however, by using query expansion, uniCOIL can achieve higher effectiveness. Following the method proposed by Zhuang and Zuccon [33] for query expansion with TILDE, we will show how TILDEv2 will act when we expand documents with TILDE. For TILDEv2, we use the implementation from the authors' code repository.<sup>4</sup>

### 2.4 Cross-encoder BERT Ranker

The state-of-the-art results on SciDocs is reported by Abolghasemi et al. [1] where they use a multi-task optimized cross-encoder BERT ranker [22]. The cross-encoder BERT ranker uses the concatenation of query and the document as the input to a BERT encoder. The BERT encoder is then followed by a projection layer  $W_{proj}$  on top of its  $[CLS]$  token to compute the relevance score:

$$s(q, d) = \text{BERT}([CLS] q [SEP] d [SEP])[CLS] * W_{proj} \quad (4)$$

In this equation,  $q$  and  $d$  represent the query and the document respectively and  $[CLS]$  as well as  $[SEP]$  are special BERT tokens [8].

<sup>3</sup><https://github.com/ielab/TILDE>

<sup>4</sup><https://github.com/ielab/TILDE/tree/main/TILDEv2>

### 3 METHODS AND EXPERIMENTAL SETTINGS

In this section, we provide details and preliminaries about our methods and experimental settings.

#### 3.1 Evaluation Benchmark

We run our experiments on the SciDocs benchmark [5]. This dataset was originally introduced as a benchmark for representation learning tasks. Later, several works including [1, 19] used the tasks of {co-view, co-read, citation, co-citation}-prediction from this benchmark as a query-by-example retrieval setting. As Figure 2 depicts, in this setting, given a query document, the goal is to retrieve and rank the most relevant documents out of a collection. The evaluation dataset for each of these four tasks includes approximately 30K total papers from a held-out pool of papers, consisting of 1K query papers and a candidate set of up to 5 positive papers and 25 negative papers [5].

To make our results comparable, we follow the prior work on SciDocs to prepare the same training data [1]. To this aim, we take the validation set of each of tasks and use 85% of them as training and 15% of them as the validation. Thus, each query in the train set has 5 relevant documents and 25 non-relevant documents. While TILDE is trained over relevant query-document pairs [34], TILDEv2 needs triplets in the format of (query, positive document, negative document). To prepare these triplets we pick two non-relevant documents per relevant document. By doing so, we create 10 triplets out of 30 training samples for each query. It should be noted that following Cohan et al. [5] we use a concatenation of abstract and title of the papers as documents.

#### 3.2 BERT-based Tokenization in Traditional Models

In order to address *RQ2*, we will examine the effects of transformer-based tokenizers as text pre-processor for traditional retrieval models. Doing so aligns the index vocabulary of traditional models with that of TILDE and TILDEv2, which in turn makes our comparison more fair. Transformers use different tokenization mechanisms e.g. WordPiece [31], which result in different query and document representations compared to common word-based tokenization approaches that are sometimes combined with normalization steps such as stemming and lemmatizing. Kamps et al. [14] show that using the BERT tokenizer as a pre-processor for BM25 results in a higher efficiency at the cost of a small decrease in effectiveness on the TREC 2020 Deep Learning Track [6]. QBE retrieval, however, has the challenge of long queries. In this work, investigate whether the same effect applies to a QBE retrieval setting. To this aim, we use the BERT<sub>base</sub> tokenizer as a pre-processor for LM and BM25.

In addition, we use the SciBERT tokenizer, which is a domain-specific BERT tokenizer, to find out if a domain-specific tokenizer would have a different effect in comparison to the BERT<sub>base</sub> tokenizer. We use three different pre-processing setups in Elasticsearch to compare with our two transformer-based tokenizers:

- Elasticsearch Standard Analyzer (SA)
- Lowercase token filter, Porter Stemmer, Whitespace tokenizer (STM1)
- Lowercase token filter, Porter Stemmer, Standard tokenizer (STM2)

In Table 2, models corresponding to these setups respectively have SA, STM1, and STM2 as their subscript. BERT-Token and SciBERT-Token as subscripts stand for using BERT and SciBERT tokenizers as the text pre-processors.

#### 3.3 Interpolation between BM25 and TILDE (TILDEv2) scores

To answer *RQ3* about the difference between BM25 and TILDE (as well as TILDEv2) in terms of their relevance signals, following Wang et al. [30], we evaluate the effect of the interpolation between the relevance scores from BM25 and from the contextualized term-based ranking models TILDE and TILDEv2. To this aim the interpolated score is computed as following:

$$s(q, d) = \alpha * s_{BM25}(q, d) + (1 - \alpha) * s_{contextualized}(q, d) \quad (5)$$

Here,  $s_{BM25}$  stands for the BM25 score for query  $q$ , and document  $d$ , and  $s_{contextualized}$  refers to the relevance score from TILDE or TILDEv2. Also,  $\alpha$  is the hyperparameter that controls the impact of the scores from BM25 and TILDE (or TILDEv2). Prior to the interpolation both of the relevance scores are normalized using z-scaling (subtracting the mean and dividing by the standard deviation). We optimize  $\alpha$  on the validation set.

Additionally, to further investigate the impact of interpolation, we do a per-query oracle interpolation in which we assume the best interpolation setting, i.e., optimal  $\alpha$ , could be predicted per query, and thus we can explore how much effectiveness is reachable by the interpolation of the scores. In the following of the paper, “oracle interpolation” refers to this latter interpolation setup and “non-oracle interpolation” refers to the vanilla interpolation, i.e., one  $\alpha$  for all queries that is optimized on the validation set.

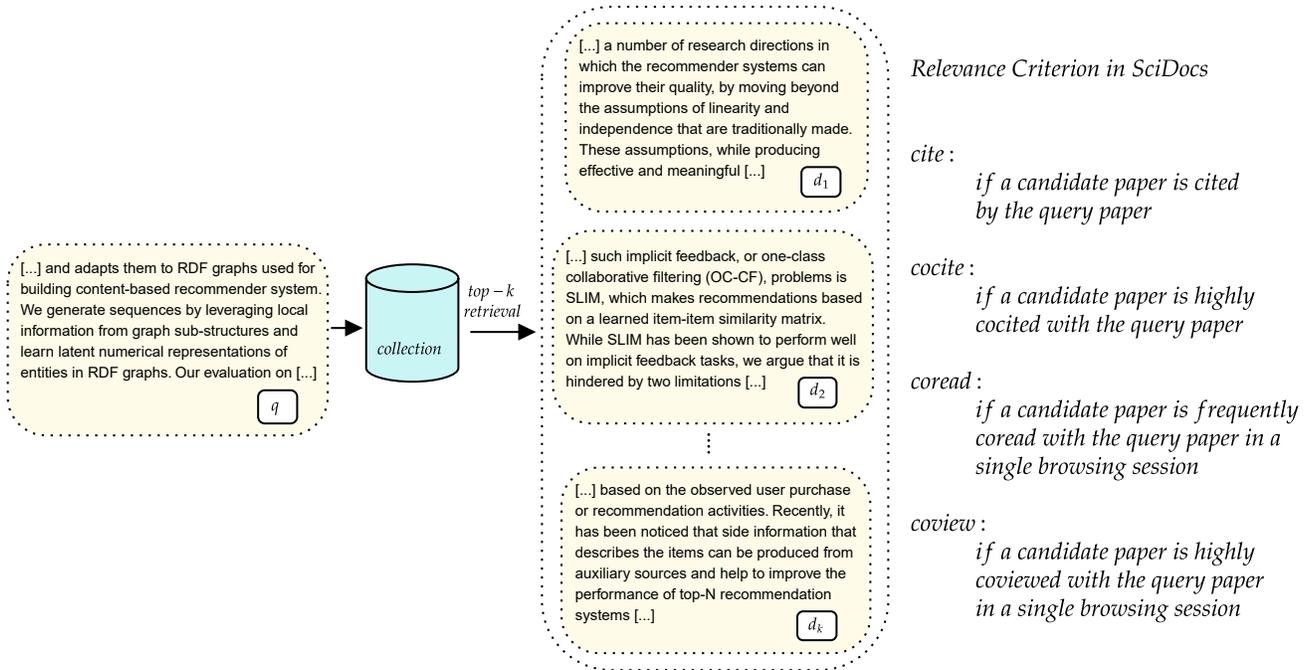
#### 3.4 Document Expansion with TILDE

The Average token count of SciDocs documents (abstract+title) is 219 and 208 for BERT and SciBERT respectively. Their 90% token count quantiles are 341 and 385. Comparing these numbers to the maximum input length of BERT models, i.e., 512 tokens, we can see a capacity for the expansion of the documents. To further investigate *RQ1*, following recent works which use document expansion to alleviate the vocabulary mismatch in contextualized term-based retrieval [16, 33], we evaluate the impact of retrieval on documents which are expanded at indexing time.

To this aim, we use TILDE in the same way as the original paper [33]. TILDE is at an advantage where it is more efficient than doc2query [23]. In this work, using TILDE<sub>SciBERT</sub> of which we found it performs the best compared to other TILDE models (Table 1), we generate  $m = 200$ , and  $m = 300$  expansion terms for TILDEv2<sub>SciBERT</sub>. It is noteworthy that similar to the original paper [33] not all expansion terms are added to a document, but only new expansion terms – that are not yet present in the document – are added.

#### 3.5 Domain-Specific BERT in TILDE and TILDEv2

To answer *RQ1*, and *RQ4*, we will investigate the power that can be brought by domain-specific pre-training to term-based ranking models. To do so, we evaluate the models’ ranking quality in



**Figure 2: In the Query-by-Example retrieval setting, given a document (in its meaning as a unit of retrieval [17]) as the query  $q$ , the goal is to retrieve and rank the top- $k$  relevant documents  $\{d_1, d_2, \dots, d_k\}$  out of a collection of documents. We use the four QBE tasks from SciDocs [5] benchmark including  $\{cite, cocite, coread, coview\}$ , each of which has its own relevance criterion [5].**

three settings: a) using  $BERT_{base}$  as encoder, b) zero-shot utilization of TILDE and TILDEv2 models which are already fine-tuned on MSMARCO, and c) using a domain-specific pre-trained BERT as their encoder. Specifically, we use SciBERT [3] since our evaluation benchmark is from the scientific domain.

### 3.6 Implementation Details

We run our experiments on NVIDIA RTX 3090 GPU machines with 24GB GPU memory. For  $BERT_{base}$ , and SciBERT we use the pre-trained models available on Huggingface. All BERT-based models are trained for 5 epochs. We use the Adam optimizer [15] with a learning rate of  $2 \times 10^{-5}$  for TILDE, and the AdamW optimizer with a learning rate of  $5 \times 10^{-6}$  for TILDEv2. In addition, we relax the maximum document length to the maximum input length of BERT during indexing.

## 4 RESULTS

**RQ1.** How effective are TILDE and TILDEv2 in query-by-example retrieval? and **RQ4** To what extent does a highly tailored domain-specific pre-trained BERT model affect the effectiveness of TILDE and TILDEv2 in comparison to when we use a  $BERT_{base}$  model?

As Table 1 shows, TILDE and TILDEv2 are less effective than a cross-encoder BERT ranker in QBE retrieval despite having longer queries. This could be due to the fact that the cross-encoder BERT ranker applies all-to-all attention across tokens in both the query and the document [17] and thus, query terms and document terms are highly contextualized for the estimation of the relevance score.

In addition, we see that TILDEv2 $BERT$  outperforms TILDE $BERT$  despite TILDEv2 being highly prone to the vocabulary mismatch problem. One hypothesis for this observation could be that in a domain-specific retrieval setup like ours, TILDEv2 with the  $BERT_{base}$  encoder predicts more effective document term weights than the term weights predicted for all tokens in the BERT vocabulary by TILDE with the  $BERT_{base}$  encoder.

In addition, using SciBERT as our domain-specific pre-trained BERT model unsurprisingly improves the ranking quality of both TILDE and TILDEv2; however, this improvement is higher between TILDE $BERT$  and TILDE $SciBERT$  than between TILDEv2 $BERT$  and TILDEv2 $SciBERT$  to an extent where TILDE $SciBERT$  even outperforms both TILDEv2 $BERT$  and TILDEv2 $SciBERT$ . This observation could be due to the fact that the vocabulary mismatch problem caused by exact matching limits the TILDEv2 ranking quality, even if we use a highly tailored domain-specific BERT as its encoder. In this respect, we investigate the impact of token-based query expansion (see section 3.4) with TILDE on the ranking quality of TILDEv2 in our QBE retrieval setting. Lines  $i$ , and  $j$  in Table 1 are the ranking results on the documents that are expanded using TILDE with the method introduced by Zhuang and Zuccon [33]. Here, we are interested to find out if using document expansion is able to compensate for the gap in the ranking quality between TILDE $SciBERT$ , and TILDEv2 $SciBERT$ .

As shown in Table 1, TILDEv2 $SciBERT$  with  $m = \{200, 300\}$  expansion terms, is still less effective than TILDE $SciBERT$ . Furthermore,

**Table 1: Ranking quality on the four SciDocs benchmark tasks using contextualized term-based ranking and cross-encoder BERT. “BERT” and “SciBERT” refers to the pre-trained model used as the encoder. “MSMARCO” indicates the utilization of TILDE or TILDev2 which are already fine-tuned on MSMARCO. Rows  $i$  and  $j$  refer to the experiments on expanded documents with  $m$  terms using TILDE<sub>SciBERT</sub> as described in section 3.4. Statistical significance improvements are according to paired t-test ( $p < 0.05$ ) with Bonferroni correction for multiple testing. Rows  $a$  and  $b$  are included from Table 2 for ease of comparison.**

Model	Co-view		Co-read		Co-cite		Cite	
	MAP	nDCG	MAP	nDCG	MAP	nDCG	MAP	nDCG
a) BM25 <sub>STM2</sub>	80.8% <sup>bcdf-j</sup>	0.9032 <sup>cdfgj</sup>	81.31% <sup>dfg</sup>	0.9112 <sup>cdg</sup>	81.53% <sup>bcdfg</sup>	0.9171 <sup>cdfg</sup>	79.74% <sup>bdg</sup>	0.9085 <sup>dg</sup>
b) BM25 <sub>SciBERT-Token</sub>	80.08% <sup>cdfg</sup>	0.8992 <sup>cdg</sup>	80.97% <sup>dfg</sup>	0.9105 <sup>dg</sup>	80.83% <sup>cdfg</sup>	0.9141 <sup>cdg</sup>	79.03% <sup>dg</sup>	0.9051 <sup>dg</sup>
c) TILDE <sub>BERT</sub>	76.74% <sup>d</sup>	0.8761 <sup>d</sup>	80.57% <sup>dg</sup>	0.8983 <sup>d</sup>	79.7% <sup>dg</sup>	0.8999 <sup>d</sup>	82.15% <sup>abdg</sup>	0.914 <sup>dg</sup>
d) TILDE <sub>MSMARCO</sub>	68.22%	0.8261	66.75%	0.8206	65.21%	0.8145	65.29%	0.8186
e) TILDE <sub>SciBERT</sub>	82.6% <sup>a-df-j</sup>	0.9115 <sup>bcdf-j</sup>	85.03% <sup>a-df-j</sup>	0.9256 <sup>a-df-j</sup>	86.38% <sup>a-df-j</sup>	0.9375 <sup>a-df-j</sup>	87.74% <sup>a-df-j</sup>	0.9431 <sup>a-dfghj</sup>
f) TILDev2 <sub>BERT</sub>	79.17% <sup>cdg</sup>	0.8948 <sup>cd</sup>	80.16% <sup>dg</sup>	0.9051 <sup>dg</sup>	80.22% <sup>dg</sup>	0.9103 <sup>dg</sup>	82.54% <sup>abdg</sup>	0.9230 <sup>abdg</sup>
g) TILDev2 <sub>MSMARCO</sub>	77.84% <sup>cd</sup>	0.8876 <sup>d</sup>	78.53% <sup>d</sup>	0.8959 <sup>d</sup>	78.17% <sup>d</sup>	0.9006 <sup>d</sup>	75.62% <sup>d</sup>	0.8866 <sup>d</sup>
h) TILDev2 <sub>SciBERT</sub>	79.59% <sup>cdg</sup>	0.8961 <sup>cdg</sup>	80.74% <sup>dg</sup>	0.9080 <sup>dg</sup>	80.94% <sup>cdfg</sup>	0.9123 <sup>dg</sup>	84.18% <sup>a-dfg</sup>	0.9314 <sup>a-dfg</sup>
TILDev2 <sub>SciBERT</sub>								
i) expansion w/ $m=200$	80.06% <sup>cdfgj</sup>	0.8985 <sup>cdg</sup>	81.29% <sup>dfgh</sup>	0.9096 <sup>dg</sup>	81.62% <sup>cdfg</sup>	0.9153 <sup>cdg</sup>	86.42% <sup>a-dfghj</sup>	0.9412 <sup>a-dfghj</sup>
j) expansion w/ $m=300$	79.38% <sup>cdg</sup>	0.8942 <sup>cd</sup>	81.17% <sup>dfg</sup>	0.9099 <sup>dg</sup>	81.93% <sup>bcdfgh</sup>	0.9165 <sup>cdg</sup>	84.4% <sup>a-dfg</sup>	0.9319 <sup>a-dfg</sup>
k) Cross-Encoder <sub>SciBERT</sub>	85.2% <sup>a-j</sup>	0.925 <sup>a-j</sup>	87.5% <sup>a-j</sup>	0.940 <sup>a-j</sup>	89.7% <sup>a-j</sup>	0.955 <sup>a-j</sup>	94.0% <sup>a-j</sup>	0.975 <sup>a-j</sup>
l) Cross-Encoder <sub>MTFT-SciBERT</sub>	86.2% <sup>a-j</sup>	0.930 <sup>a-j</sup>	87.7% <sup>a-j</sup>	0.940 <sup>a-j</sup>	91.0% <sup>a-j</sup>	0.961 <sup>a-j</sup>	94.2% <sup>a-j</sup>	0.976 <sup>a-j</sup>

**Table 2: Ranking quality of traditional retrieval models on the four SciDocs benchmark tasks with different tokenization approaches. SA, STM1, STM2, BERT-Token, and SciBERT-Token refer to the pre-processing setting as described in section 3.2. Statistical significance improvements are according to paired t-test ( $p < 0.05$ ) with Bonferroni correction for multiple testing.**

Model	Co-view		Co-read		Co-cite		Cite	
	MAP	nDCG	MAP	nDCG	MAP	nDCG	MAP	nDCG
a) LM <sub>SA</sub>	74.78%	0.8724	74.32% <sup>b</sup>	0.8750	74.64%	0.8812	71.30%	0.8653
b) LM <sub>STM1</sub>	74.82%	0.8737	73.51%	0.8694	74.60%	0.8810	70.98%	0.8636
c) LM <sub>STM2</sub>	75.74% <sup>abde</sup>	0.8786 <sup>abe</sup>	74.90% <sup>ab</sup>	0.8771 <sup>b</sup>	75.80% <sup>abde</sup>	0.8873 <sup>ab</sup>	72.15% <sup>abe</sup>	0.8696 <sup>b</sup>
d) LM <sub>BERT-Token</sub>	74.9%	0.8734	74.76% <sup>b</sup>	0.8778 <sup>b</sup>	74.95%	0.8829	72.04% <sup>abe</sup>	0.8694
e) LM <sub>SciBERT-Token</sub>	74.74%	0.8717	74.69% <sup>b</sup>	0.8771 <sup>b</sup>	74.81%	0.8827	71.46%	0.8666
f) BM25 <sub>SA</sub>	77.86% <sup>a-e</sup>	0.8876 <sup>a-e</sup>	78.03% <sup>a-e</sup>	0.8949 <sup>a-e</sup>	77.95% <sup>a-e</sup>	0.8994 <sup>a-e</sup>	76.12% <sup>a-e</sup>	0.8892 <sup>a-e</sup>
g) BM25 <sub>STM1</sub>	80.21% <sup>a-f</sup>	0.9002 <sup>a-f</sup>	80.52% <sup>a-f</sup>	0.9074 <sup>a-f</sup>	80.85% <sup>a-f</sup>	0.9137 <sup>a-f</sup>	79.03% <sup>a-f</sup>	0.9048 <sup>a-f</sup>
h) BM25 <sub>STM2</sub>	80.8% <sup>a-gij</sup>	0.9032 <sup>a-gi</sup>	81.31% <sup>a-gi</sup>	0.9112 <sup>a-g</sup>	81.53% <sup>a-gij</sup>	0.9171 <sup>a-g</sup>	79.74% <sup>a-gij</sup>	0.9085 <sup>a-g</sup>
i) BM25 <sub>BERT-Token</sub>	79.76% <sup>a-f</sup>	0.8974 <sup>a-f</sup>	80.61% <sup>a-f</sup>	0.9088 <sup>a-f</sup>	80.5% <sup>a-f</sup>	0.9125 <sup>a-f</sup>	79.19% <sup>a-f</sup>	0.9057 <sup>a-f</sup>
j) BM25 <sub>SciBERT-Token</sub>	80.08% <sup>a-f</sup>	0.8992 <sup>a-f</sup>	80.97% <sup>a-gi</sup>	0.9105 <sup>a-f</sup>	80.83% <sup>a-fi</sup>	0.9141 <sup>a-f</sup>	79.03% <sup>a-f</sup>	0.9051 <sup>a-f</sup>

the table shows that the ranking quality of BM25<sub>STM2</sub> on the original documents (line  $a$ ) is still comparable with the ranking quality of TILDev2<sub>SciBERT</sub> on the expanded documents (lines  $i$  and  $j$ ). It is noteworthy that to make sure we are expanding the documents with enough tokens we investigate the average number of tokens added to the documents by the expansion with TILDE<sub>SciBERT</sub>. By doing so, we find that for  $m = 200$ , and  $m = 300$ , approximately 49 and 128 new tokens are appended to the documents on average. Additionally, we find that using  $m = 100$  results in roughly 2.6

new tokens on average. These numbers beside the statistics of the tokens in SciDocs benchmark, provided in Section 3.4, indicate that  $m$  should be tuned in order to take advantage from the document expansion with TILDE in QBE retrieval setting. Finally, we see that the zero-shot utilization of TILDE<sub>MSMARCO</sub> and TILDev2<sub>MSMARCO</sub> does not show superior performance over the fine-tuned TILDE and TILDev2 with both BERT and SciBERT encoders. It should be noted that taking models which are already fine-tuned on general domain (like TILDE<sub>MSMARCO</sub> and TILDev2<sub>MSMARCO</sub>) and further

**Table 3: Results for non-oracle interpolation (the interpolation parameter  $\alpha$  is optimized on the validation set) between  $BM25_{STM2}$ ,  $TILDE_{SciBERT}$ , and  $TILDEV2_{SciBERT}$ . Statistical significance improvements are according to paired t-test ( $p < 0.05$ ) with Bonferroni correction for multiple testing. Rows *a*, *b*, and *c* are included from Table 1 for ease of comparison.**

Model	Co-view		Co-read		Co-cite		Cite	
	MAP	nDCG	MAP	nDCG	MAP	nDCG	MAP	nDCG
a) $BM25_{STM2}$	80.8% <sup>c</sup>	0.9032 <sup>c</sup>	81.31%	0.9112	81.53%	0.9171	79.74%	0.9085
b) $TILDE_{SciBERT}$	82.6% <sup>ac</sup>	0.9115 <sup>c</sup>	85.03% <sup>ace</sup>	0.9256 <sup>ac</sup>	86.38% <sup>ace</sup>	0.9375 <sup>ace</sup>	87.74% <sup>ace</sup>	0.9431 <sup>ace</sup>
c) $TILDEV2_{SciBERT}$	79.59%	0.8961	80.74%	0.9080	80.94%	0.9123	84.18% <sup>a</sup>	0.9314 <sup>a</sup>
d) $BM25_{STM2} + TILDE_{SciBERT}$	<b>85.29%<sup>abce</sup></b>	<b>0.9214<sup>abce</sup></b>	<b>86.52%<sup>abce</sup></b>	<b>0.9395<sup>abce</sup></b>	<b>88.32%<sup>abce</sup></b>	<b>0.9494<sup>abce</sup></b>	<b>88.46%<sup>abce</sup></b>	<b>0.9496<sup>abce</sup></b>
e) $BM25_{STM2} + TILDEV2_{SciBERT}$	81.56% <sup>ac</sup>	0.9032 <sup>c</sup>	82.63% <sup>ac</sup>	0.9183 <sup>ac</sup>	83.06% <sup>ac</sup>	0.9242 <sup>ac</sup>	84.18% <sup>a</sup>	0.9318 <sup>a</sup>

fine-tuning them on the task domain is a typical approach which could result in improvement in their ranking quality; however, we leave this item as a direction to be explored in future work.

**RQ2.** *What is the effectiveness of traditional lexical matching models with varying tokenization strategies in comparison to  $TILDE$  and  $TILDEV2$*

Table 2 shows that leveraging BERT and SciBERT tokenizers results in competitive ranking quality in both probabilistic language model based retrieval and BM25 in comparison to the three traditional pre-processing setups introduced in section 3.2.

Moreover, as the results of Table 1 shows, the ranking quality of  $BM25_{STM2}$  not only outperforms LM and BM25 with different traditional and BERT-based pre-processing approaches, but also it could even outperform  $TILDE_{BERT}$ , and  $TILDEV2_{BERT}$  in most of the tasks. In fact, we do not see a large gap between BM25 compared to  $TILDEV2$  as was shown for retrieval based on short queries in the experiments on MSMARCO and TREC DL Track benchmarks [33]. This finding is important as (1) it sheds light on the challenges of retrieval settings different from the common evaluation benchmarks including MSMARCO and the TREC DL Track; (2) raises the question how effective other contextualized term-based ranking models would be in those settings.

**RQ3.** *To what extent do  $TILDE$  and  $TILDEV2$  encode a different relevance signal from  $BM25$ ?*

The blue lines in Figure 3 show the ranking quality for  $TILDE_{SciBERT}$  and  $TILDEV2_{SciBERT}$  when their scores are interpolated with the BM25 score over varying values of interpolation parameter  $\alpha$  with the step of 0.1. Besides, Table 3 shows the ranking quality for the interpolations with the  $\alpha$  that is tuned over the validation set. We can see that an optimal interpolation between the scores from BM25 and the contextualized term-based ranking models  $TILDE$  and  $TILDEV2$  could provide significant improvements for almost all tasks over the individual rankers participating in the interpolation. The only exceptions are in the *co-view*, and *cite* tasks. To be specific, there is no improvement over BM25 in the nDCG metric in the *co-view* (line e vs. line a in Table 3). Besides, in the *cite* task the improvement over  $TILDEV2$  (line e vs. line c in Table 3) is not significant for the nDCG metric, and there is no improvement for the MAP metric. Nevertheless, the improvements obtained by the interpolation for almost all tasks and metrics indicates that  $TILDE$  and  $TILDEV2$  are capturing different relevance signals compared to BM25.

To further investigate the impact of the score interpolation with BM25 scores, we perform an oracle interpolation in which we assume the optimal interpolation hyperparameter  $\alpha$  is known for each individual query. This query-specific optimal value is selected over varying values of  $\alpha$  with the step of 0.1. Table 4 as well as orange lines in Figure 3 show the results for the oracle interpolation. We can see that the oracle interpolation would result in a substantial improvement for both  $TILDE$  and  $TILDEV2$ .

Moreover, we can see in Table 4 that there is a subset of queries for which the BM25 ranking alone is better than the interpolation (queries with optimal  $\alpha=1$ ). This number is lower for the interpolation with  $TILDE$  than for the interpolation with  $TILDEV2$ . One hypothesis for this observation could be that the interpolation with  $TILDE$  is likely to be more helpful for BM25 since  $TILDE$  could bring more contextualization power for BM25 as it incorporates the term importance for all tokens in the query. In other words, since  $TILDEV2$  pre-computes term weights only for the tokens of the document (whereas  $TILDE$  pre-computes the term importance weight for all the tokens in the BERT vocabulary per document), due to the chance of vocabulary mismatch in  $TILDEV2$ , it could incorporate less query-dependent contextualization than  $TILDE$ .

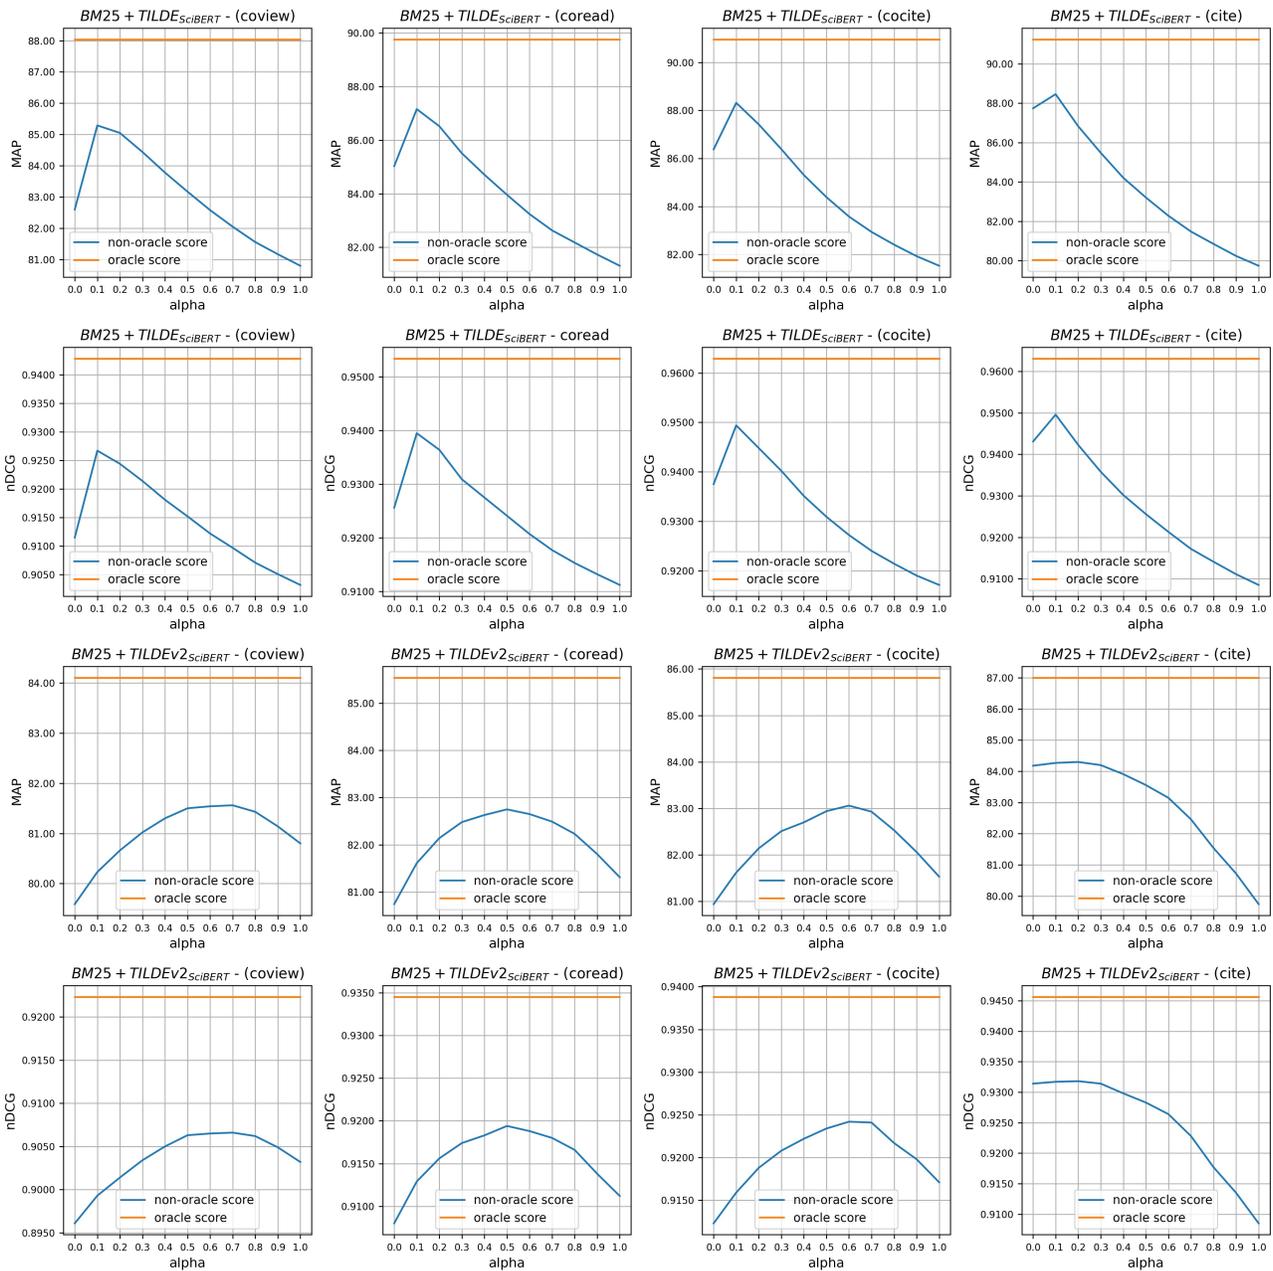
In addition, we see that the margin between the oracle interpolation results and both non-interpolated scores as well as non-oracle interpolation scores (Table 3) is substantial, which demonstrates that more complex aggregation methods could benefit more from the relevance signals from  $TILDE$ ,  $TILDEV2$  and BM25.

## 5 DISCUSSION

In this section, we further analyze the interpolation between BM25 and  $TILDE$  ( $TILDEV2$ ) in terms of the interpolation effectiveness and the interpolation weight  $\alpha$ .

### 5.1 Interpolation effectiveness

The first two rows on the top of Figure 3 correspond to the interpolation between  $TILDE$  and BM25 and the two rows in the bottom correspond to the interpolation between  $TILDEV2$  and BM25. Comparing the nDCG and MAP plots for the interpolation between  $TILDE$  and BM25, we can see that for this combination,  $\alpha=0.1$  shows the highest ranking quality for both nDCG and MAP metrics in all tasks. Thus, a high weight for  $TILDE$  with a small weight for BM25



**Figure 3: Results for TILDE and TILDEv2 with varying values of interpolation parameter  $\alpha$ . The lines in blue and orange represent the effectiveness based on the non-oracle and oracle interpolations respectively.  $\alpha = 0.0$  represents the TILDE- or TILDEv2-only setting;  $\alpha = 1.0$  represents the BM25-only setting.**

gives the highest effectiveness for this combination. This observation could mean that while TILDE, as contextualized transformer-based model, is able to outperform BM25 as an exact matching model, it could still benefit from the strong lexical relevance scores from BM25.

On the other hand, for the combination of TILDEv2 and BM25 we see that the highest ranking quality is obtained with  $\alpha \in \{0.3,$

$0.4, 0.5, 0.6, 0.7, 0.8\}$  depending on the task. The exceptions are in the *cite*, and *coview* tasks as described in the answer to *RQ3* in Section 4. Thus, in the combination of TILDEv2 and BM25, an equal or slightly higher weight for BM25 relative to TILDEv2 gives the optimal results. A hypothesis for this observation could be that while both BM25 and TILDEv2 are performing based on exact matching, the term weights from TILDEv2, which are predicted

**Table 4: Results for oracle interpolation (optimal  $\alpha$  per query) between  $BM25_{STM2}$ ,  $TILDE_{SciBERT}$ , and  $TILDEv2_{SciBERT}$ . Statistical significance with paired t-test ( $p < 0.05$ ) is reported only with respect to non-interpolated scores ( $\alpha = 0$ ) of these three models in Table 1 (row a, e and h).  $\alpha_{average}$  represents the mean of the optimal  $\alpha$  values picked per query. The number of queries with optimal  $\alpha = 0$  stands for the number of queries for which the interpolation does not improve their effectiveness compared to TILDE or TILDEv2 only.**

Model	Co-view		Co-read		Co-cite		Cite	
	MAP	nDCG	MAP	nDCG	MAP	nDCG	MAP	nDCG
$BM25_{+oracle} TILDE_{SciBERT}$								
ranking quality	88.04% <sup>ae</sup>	0.9428 <sup>ae</sup>	89.76% <sup>ae</sup>	0.9534 <sup>ae</sup>	90.96% <sup>ae</sup>	0.9629 <sup>ae</sup>	91.24% <sup>ae</sup>	0.9631 <sup>ae</sup>
$\alpha_{average}$	0.1265	0.1294	0.1048	0.1053	0.1044	0.1048	0.0839	0.0857
#queries with optimal $\alpha = 0$	537	533	563	557	550	552	624	619
#queries with optimal $\alpha = 1$	19	21	6	5	14	13	4	4
IQR of the optimal $\alpha$ over queries	0.1	0.2	0.1	0.1	0.1	0.1	0.1	0.1
$BM25_{+oracle} TILDEv2_{SciBERT}$								
ranking quality	84.10% <sup>ah</sup>	0.9223 <sup>ah</sup>	85.54% <sup>ah</sup>	0.9345 <sup>ah</sup>	85.81% <sup>ah</sup>	0.9388 <sup>ah</sup>	87.00% <sup>ah</sup>	0.9456 <sup>ah</sup>
$\alpha_{average}$	0.3169	0.3205	0.3040	0.3048	0.3337	0.3339	0.2073	0.2083
#queries with optimal $\alpha = 0$	467	463	446	447	419	420	575	573
#queries with optimal $\alpha = 1$	84	84	71	72	92	94	33	33
IQR of the optimal $\alpha$ over queries	0.7	0.7	0.6	0.6	0.6	0.6	0.4	0.4

through contextualization of the document terms, are not always more effective than the term scores from BM25; however, they can act as a complement for each other and thus their interpolation could benefit from both.

## 5.2 Interpolation weight

To further analyze the interpolation weight  $\alpha$ , we consider the two aforementioned settings of oracle interpolation and non-oracle interpolation.

*Non-oracle interpolation.* We can see in Figure 3 (blue lines) that for the effective interpolations, i.e, the interpolations that result in higher effectiveness than each individual ranker included in the interpolation, the interpolation weight  $\alpha$  in the combination of BM25 and TILDEv2 has a wider range than in the combination of BM25 and TILDE. This indicates that in this experimental setting the interpolation of BM25 and TILDEv2 could be achieved by a broader range of  $\alpha$  values and is therefore more robust to the choice of interpolation weight than for BM25 and TILDE.

*Oracle interpolation.* As a measure of the statistical dispersion, we report the inter-quartile range (IQR) for the oracle interpolation weight  $\alpha$  which is shown in Table 4. Taking the range of  $\alpha$  [0.0, 1.0] into account, we can see that we have low inter-quartile range (IQR) for the optimal values of  $\alpha$  per query in the interpolation with TILDE (top part of the table). On the other hand, the IQR for the optimal values of  $\alpha$  per query for the interpolation with TILDEv2 are much higher (bottom part of the table), which indicates that the optimal interpolation setting for the queries are more varied. This observation could give some sense of robustness against query variation for TILDE in comparison to TILDEv2 in this experimental setting. In other words, a query-dependent approach for optimizing  $\alpha$  would be more robust against query variation for TILDE than for TILDEv2.

## 6 CONCLUSION

In this paper we investigated the generalizability of two contextualized term-based ranking models TILDE and TILDEv2 for a QBE retrieval setting. In QBE, the queries are much longer than in ad-hoc retrieval, and efficient query processing is essential. We were specifically interested to see to what extent the relative performance of contextualized term-based ranking models in comparison to both traditional term-based models and the effective cross-encoder BERT ranker is generalizable to a QBE retrieval setting.

Our results show that similar to the original papers [33, 34], TILDE and TILDEv2 are less effective than a cross-encoder BERT ranker in QBE retrieval despite the context of longer queries. On the other hand, in the original papers, TILDE and TILDEv2 have shown superior ranking quality in comparison to BM25 as a traditional term-based retrieval model. We investigated if the same pattern exists in a query-by-example retrieval setting and our results show that BM25 has a competitive ranking quality compared to TILDE and TILDEv2. In fact, not only is it competitive, but also in some cases it could outperform TILDE and TILDEv2.

This finding is important as (1) it sheds light on the challenges of retrieval settings different from the common evaluation benchmarks including MSMARCO and the TREC DL Track; (2) raises the question how effective other contextualized term-based ranking models would be in those settings. Our results indicate that QBE retrieval is structurally different from other IR settings and requires special attention for methods development.

Furthermore, we investigated the impact of the interpolation between BM25 and TILDE as well as TILDEv2. By doing so, we find that a linear interpolation between the score of TILDE (TILDEv2) with that of BM25 leads to an improvement in the ranking effectiveness. This shows that the relevance signals from contextualized ranking models TILDE and TILDEv2 are complementary to the relevance signals from BM25. Additionally, through an analysis on the oracle interpolation between BM25 and TILDE (TILDEv2), we

show that more stratified approaches could benefit more from the interpolation between the scores from these models.

## 7 ACKNOWLEDGMENTS

This work is funded by the DoSSIER project under European Union’s Horizon 2020 research and innovation program, Marie Skłodowska-Curie grant agreement No. 860721.

## REFERENCES

- [1] Amin Abolghasemi, Suzan Verberne, and Leif Azzopardi. 2022. Improving BERT-based Query-by-Document Retrieval with Multi-Task Optimization. In *Advances in Information Retrieval, 44th European Conference on IR Research, ECIR 2022*. <https://arxiv.org/abs/2202.00373>
- [2] A Askari and S Verberne. 2021. Combining lexical and neural retrieval with longformer-based summarization for effective case law retrieval. In *Proceedings of the second international conference on design of experimental search & information REtrieval systems*. CEUR, 162–170.
- [3] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. 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)*. 3615–3620.
- [4] Adam Berger and John Lafferty. 2017. Information retrieval as statistical translation. In *ACM SIGIR Forum*, Vol. 51. ACM New York, NY, USA, 219–226.
- [5] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 2270–2282. <https://doi.org/10.18653/v1/2020.acl-main.207>
- [6] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. Overview of the TREC 2020 deep learning track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference. NIST Special Publication*.
- [7] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, Ellen M Voorhees, and Ian Soboroff. 2021. TREC deep learning track: Reusable test collections in the large data regime. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2369–2375.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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)*. 4171–4186.
- [9] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE v2: Sparse lexical and expansion model for information retrieval. *arXiv preprint arXiv:2109.10086* (2021).
- [10] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE: Sparse lexical and expansion model for first stage ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2288–2292.
- [11] Atsushi Fujii, Makoto Iwayama, and Noriko Kando. 2007. Overview of the Patent Retrieval Task at the NTCIR-6 Workshop. In *NTCIR*.
- [12] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2021. COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 3030–3042.
- [13] Djoerd Hiemstra. 1998. A linguistically motivated probabilistic model of information retrieval. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 569–584.
- [14] Jaap Kamps, Nikolaos Kondylidis, and David Rau. 2020. Impact of Tokenization, Pretraining Task, and Transformer Depth on Text Ranking. In *TREC*.
- [15] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- [16] Jimmy Lin and Xueguang Ma. 2021. A few brief notes on deepimpact, coil, and a conceptual framework for information retrieval techniques. *arXiv preprint arXiv:2106.14807* (2021).
- [17] Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. 2021. Pretrained transformers for text ranking: Bert and beyond. *Synthesis Lectures on Human Language Technologies* 14, 4 (2021), 1–325.
- [18] Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonellotto. 2021. *Learning Passage Impacts for Inverted Indexes*. Association for Computing Machinery, New York, NY, USA, 1723–1727. <https://doi.org/10.1145/3404835.3463030>
- [19] Sheshera Mysore, Arman Cohan, and Tom Hope. 2021. Multi-Vector Models with Textual Guidance for Fine-Grained Scientific Document Similarity. *arXiv preprint arXiv:2111.08366* (2021).
- [20] Sheshera Mysore, Tim O’Gorman, Andrew McCallum, and Hamed Zamani. 2021. CSFCube-A Test Collection of Computer Science Research Articles for Faceted Query by Example. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- [21] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. In *CoCo@ NIPS*.
- [22] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085* (2019).
- [23] Rodrigo Nogueira, Jimmy Lin, and AI Epistemic. [n.d.]. From doc2query to docTTTTTquery. ([n. d.]).
- [24] Florina Piroi and Allan Hanbury. 2019. Multilingual Patent Text Retrieval Evaluation: CLEF-IP. In *Information Retrieval Evaluation in a Changing World*. Springer, 365–387.
- [25] Florina Piroi, Mihai Lupu, Allan Hanbury, and Veronika Zenz. 2011. CLEF-IP 2011: Retrieval in the Intellectual Property Domain. In *CLEF (notebook papers/labs/workshop)*. Citeseer.
- [26] Jay M Ponte and W Bruce Croft. 2017. A language modeling approach to information retrieval. In *ACM SIGIR Forum*, Vol. 51. ACM New York, NY, USA, 202–208.
- [27] SE Robertson, Steve Walker, S Jones, and M Hancock-Beaulieu. [n.d.]. M. Gatford (1995). Okapi at trec-3. In *Proceedings of the Third Text REtrieval Conference (TREC-3)*.
- [28] Guilherme Moraes Rosa, Ruan Chaves Rodrigues, Roberto Lotufo, and Rodrigo Nogueira. 2021. Yes, bm25 is a strong baseline for legal case retrieval. *arXiv preprint arXiv:2105.05686* (2021).
- [29] Sheikh Muhammad Sarwar and James Allan. 2020. Query by Example for Cross-Lingual Event Retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1601–1604.
- [30] Shuai Wang, Shengyao Zhuang, and Guido Zuccon. 2021. Bert-based dense retrievers require interpolation with bm25 for effective passage retrieval. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*. 317–324.
- [31] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [32] Chengxiang Zhai and John Lafferty. 2004. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems (TOIS)* 22, 2 (2004), 179–214.
- [33] Shengyao Zhuang and Guido Zuccon. 2021. Fast passage re-ranking with contextualized exact term matching and efficient passage expansion. *arXiv preprint arXiv:2108.08513* (2021).
- [34] Shengyao Zhuang and Guido Zuccon. 2021. TILDE: Term Independent Likelihood model for Passage Re-ranking. (2021).