



Exploring Classic and Neural Lexical Translation Models for Information Retrieval: Interpretability, Effectiveness, and Efficiency Benefits

Leonid Boytsov^{1(✉)} and Zico Kolter^{1,2}

¹ Bosch Center for Artificial Intelligence, Pittsburgh, USA
leonid.boytsov@us.bosch.com

² Carnegie Mellon University, Pittsburgh, USA
zkolter@cs.cmu.edu

Abstract. We study the utility of the lexical translation model (IBM Model 1) for English text retrieval, in particular, its neural variants that are trained end-to-end. We use the neural Model1 as an aggregator layer applied to context-free or contextualized query/document embeddings. This new approach to design a neural ranking system has benefits for effectiveness, efficiency, and interpretability. Specifically, we show that adding an *interpretable* neural Model 1 layer on top of BERT-based contextualized embeddings (1) does not decrease accuracy and/or efficiency; and (2) may overcome the limitation on the maximum sequence length of existing BERT models. The context-free neural Model 1 is less effective than a BERT-based ranking model, but it can run efficiently on a CPU (without expensive index-time precomputation or query-time operations on large tensors). Using Model 1 we produced best neural *and* non-neural runs on the MS MARCO document ranking leaderboard in late 2020.

1 Introduction

A typical text retrieval system relies on simple term-matching techniques to generate an initial list of candidates, which can be further re-ranked using a learned model [10,13]. Thus, retrieval performance is adversely affected by a mismatch between query and document terms, which is known as a vocabulary gap problem [18,74]. Two decades ago Berger and Lafferty [4] proposed to reduce the vocabulary gap and, thus, to improve retrieval effectiveness with a help of a lexical translation model called IBM Model 1 (henceforth, simply Model 1). Model 1 has strong performance when applied to finding answers in English question-answer (QA) archives using questions as queries [35,57,65,71] as well as to cross-lingual retrieval [38,73]. Yet, little is known about its effectiveness on *realistic monolingual* English queries, partly, because training Model 1 requires large query sets, which previously were not publicly available.

Research Question 1. In the past, Model 1 was trained on question-document pairs of similar lengths which simplifies the task of finding useful associations

between query terms and terms in relevant documents. It is not clear if Model 1 can be successfully trained if queries are substantially, e.g., two orders of magnitude, shorter than corresponding relevant documents.

Research Question 2. Furthermore, Model 1 was trained in a *translation* task using an expectation-maximization (EM) algorithm [9, 16] that produces a sparse matrix of conditional translation probabilities, i.e., a *non-parametric* model. Can we do better by *parameterizing* conditional translation probabilities with a neural network and learning the model *end-to-end* in a *ranking*—rather than a translation—task?

To answer these research questions we experiment with lexical translation models on two recent MS MARCO collections, which have hundreds of thousands of real user queries [12, 49]. Specifically, we consider a novel class of ranking models where an *interpretable* neural Model 1 layer *aggregates* an output of a token-embedding neural network. The resulting composite network (including token embeddings) is learned end-to-end using a ranking objective. We consider two scenarios: context-independent token embeddings [11, 22] and contextualized token embeddings generated by BERT [17]. Note that our approach is *generic* and can be applied to other embedding networks as well.

The neural Model 1 layer produces all pairwise similarities $T(q|d)$ for all query and documents BERT word pieces, which are combined via a straightforward product-of-sum formula without any learned weights:

$$P(Q|D) = \prod_{q \in Q} \sum_{d \in D} T(q|d)P(d|D), \quad (1)$$

where $P(d|D)$ is a maximum-likelihood estimate of the occurrence of d in D . Indeed, a query-document score is a product of scores for individual query word pieces, which makes it easy to pinpoint word pieces with largest contributions. Likewise, for every query word piece we can easily identify document word pieces with highest contributions to its score. This makes our model *more interpretable* compared to prior work.

Our contributions can be summarized as follows:

1. Adding an *interpretable* neural Model 1 layer on top of BERT entails virtually no loss in accuracy *and* efficiency compared to the vanilla BERT ranker, which is not readily interpretable.
2. In fact, for long documents the BERT-based Model 1 may outperform baseline models applied to truncated documents, thus, overcoming the limitation on the maximum sequence length of existing pretrained Transformer [67] models. However, evidence was somewhat inconclusive and we found it was also not conclusive for previously proposed CEDR [44] models that too incorporate an aggregator layer (though a *non-interpretable* one);
3. A fusion of the non-parametric Model 1 with BM25 scores can outperform the baseline models, though the gain is modest ($\approx 3\%$). In contrast, the fusion with the context-free neural Model 1 can be substantially ($\approx 10\%$) more effective than the fusion with its non-parametric variant. We show that the neural

Model 1 can be sparsified and executed on a CPU more than 10^3 times faster than a BERT-based ranker on a GPU. We can, thus, improve the first retrieval stage *without expensive* index-time precomputation approaches.

2 Related Work

Translation Models for Text Retrieval. This line of work begins with an influential paper by Berger and Lafferty [4] who first applied Model 1 to text retrieval [4]. It was later proved to be useful for finding answers in *monolingual* QA archives [35, 57, 65, 71] as well as for cross-lingual document retrieval [38, 73]. Model 1 is a *non-parametric* and *lexical* translation model that learns *context-independent* translation probabilities of lexemes (or tokens) from a set of paired documents called a *parallel corpus* or *bixtext*. The learning method is a variant of the expectation-maximization (EM) algorithm [9, 16].

A generic approach to improve performance of non-parametric statistical learning models consists in parameterizing respective probabilities using neural networks. An early successful implementation of this idea in language processing were the hybrid HMM-DNN/RNN systems for speech recognition [5, 26]. More concretely, our proposal to use the neural Model 1 as a last network layer was inspired by the LSTM-CRF [32] and CEDR [44] architectures.

There is prior history of applying the neural Model 1 to retrieval, however, without training the model on a ranking task. Zuccon et al. [75] computed translation probabilities using the cosine similarity between word embeddings (normalized over the sum of similarities for top- k closest words). They achieved modest 3–7% gains on four *small-scale* TREC collections. Ganguly et al. [19] used a nearly identical approach (on similar TREC collections) and reported slightly better (6–12%) gains. Neither Zuccon et al. [75] nor Ganguly et al. [19] attempted to learn translation probabilities from a large set of real user queries.

Zbib et al. [73] employed a *context-dependent* lexical neural translation model for *cross-lingual* retrieval. They first learn context-dependent translation probabilities from a bilingual parallel corpus in a lexical *translation* task. Given a document, highest translation probabilities together with respective tokens are precomputed in advance and stored in the index. Zbib et al. [73] trained their model on aligned sentences of similar lengths. In the case of *monolingual retrieval*, however, we do not have such fine-grained training data as queries are paired only with much longer relevant documents. To our knowledge, there is no *reliable* way to obtain sentence-level relevance labels from this data.

Neural Ranking models have been a popular topic in recent years [24], but the success of early approaches—which predate BERT—was controversial [40]. This changed with adoption of large pretrained models [55], especially after the introduction of the Transformer models [17] and release of BERT [17]. Nogueira and Cho were first to apply BERT to ranking of text documents [50]. In the TREC 2019 deep learning track [12] as well as on the MS MARCO leaderboard [1], BERT-based models outperformed all other approaches by a large margin.

The Transformer model [67] uses an attention mechanism [3] where each sequence position can attend to all the positions in the previous layer.

Because self-attention complexity is quadratic with respect to a sequence length, Transformer models (BERT including) support only limited-length inputs. A number of proposals—see Tay et al. [66] for a survey—aim to mitigate this constraint, which is complementary to our work.

To process longer documents with existing pretrained models, one has to split documents into several chunks, process each chunk separately, and aggregate results, e.g., by computing a maximum or a weighted prediction score [15, 72]. Such models cannot be trained end-to-end on full documents. Furthermore, a training procedure has to assume that each chunk in a relevant document is relevant as well, which is not quite accurate. To improve upon simple aggregation approaches, MacAvaney et al. [44] combined output of several document chunks using three simpler models: KNRM [70], PACRR [33], and DRMM [23]. A more recent PARADE architectures use even simpler aggregation approaches [39]. However, none of the mentioned aggregator models is interpretable and we propose to replace them with our neural Model 1 layer.

Interpretability and Explainability of statistical models has become a busy area of research. However, a vast majority of approaches rely on training a separate explanation model or exploiting saliency/attention maps [41, 59]. This is problematic, because explanations provided by extraneous models cannot be verified and, thus, trusted [59]. Moreover, saliency/attention maps reveal which data parts are being processed by a model, but not *how* the model processes them [34, 59, 62]. Instead of producing unreliable post hoc explanations, Rudin [59] advocates for networks whose computation is transparent *by design*. If full transparency is not feasible, there is still a benefit of *last-layer* interpretability.

In text retrieval we know only two implementations of this idea. Hofstätter et al. [29] use a kernel-based formula by Xiong et al. [70] to compute soft-match counts over contextualized embeddings. Because each pair of query-document tokens produces several soft-match values corresponding to different thresholds, it is problematic to aggregate these values in an explainable way. Though this approach does offer insights into model decisions, the aggregation formula is a relatively complicated two-layer neural network with a non-linear (logarithm) activation function after the first layer [29]. ColBERT in the re-ranking mode can be seen as an interpretable interaction layer, however, unlike the neural Model 1 its use entails a 3% degradation in accuracy [37].

Efficiency. It is possible to speed-up ranking by deferring some computation to index time. They can be divided into two groups. First, it is possible to precompute separate query and document representations, which can be quickly combined at query-time in a non-linear fashion [20, 37]. This method entails little to no performance degradation. Second, one can generate (or enhance) independent query and document representations to compare them via the inner-product computation. Representations—either dense or sparse—were shown to improve the first-stage retrieval albeit at the cost of expensive indexing processing and some loss in effectiveness. In particular, Khattab et al. [36] show that *dense* representations are inferior to the vanilla BERT ranker [52] in a QA task.

In the case of sparse representations, one can rely on Transformer [67] models to generate importance weights for document or query terms [14], augment documents with most likely query terms [51, 52], or use a combination of these methods [43]. Due to sparsity of data generated by term expansion and re-weighting models, it can be stored in a traditional inverted file to improve performance of the first retrieval stage. However, these models are less effective than the vanilla BERT ranker [52] and they require *costly* index-time processing.

3 Methods

Token Embeddings and Transformers. We assume that an input text is split into small chunks of texts called *tokens*. A token can be a complete English word, a word piece, or a lexeme (a lemma). The length of a document d —denoted as $|d|$ —is measured in the number of tokens. Because neural networks cannot operate directly on text, a sequence of tokens $t_1 t_2 \dots t_n$ is first converted to a sequences of d -dimensional embedding vectors $w_1 w_2 \dots w_n$ by an *embedding* network. Initially, embedding networks were context independent, i.e., each token was always mapped to the same vector [11, 22, 46]. Peters et al. [55] demonstrated superiority of *contextualized*, i.e., context-dependent, embeddings produced a multi-layer bi-directional LSTM [21, 27, 61] pretrained on a large corpus in a *self-supervised* manner. These were later outstripped by large pretrained Transformers [17, 56].

In our work we use two types of embeddings: vanilla context-free embeddings (see [22] for an excellent introduction) and BERT-based contextualized embeddings [17]. Due to space constraints, we do not discuss BERT architecture in detail (see [17, 60] instead). It is crucial, however, to know the following:

- Contextualized token embeddings are vectors of the last-layer hidden state;
- BERT operates on word pieces [69] rather than complete words;
- The vocabulary has close to 30K tokens and includes two special tokens: [CLS] (an aggregator) and [SEP] (a separator);
- [CLS] is always prepended to every token sequence and its embedding is used as a sequence representation for classification and ranking tasks.

The “vanilla” BERT ranker uses a single fully-connected layer as a prediction head, which converts the [CLS] vector into a scalar. It makes a prediction based on the following sequence of tokens: [CLS] q [SEP] d [SEP], where q is a query and $d = t_1 t_2 \dots t_n$ is a document. Long documents and queries need to be truncated so that the overall number of tokens does not exceed 512. To overcome this limitation, MacAvaney et al. [44] proposed an approach that:

- splits longer documents d into m chunks: $d = d_1 d_2 \dots d_m$;
- generates m token sequences [CLS] q [SEP] d_i [SEP];
- processes each sequence with BERT to generate contextualized embeddings for regular tokens as well as for [CLS].

The outcome of this procedure is m [CLS]-vectors cls_i and n contextualized vectors $w_1 w_2 \dots w_n$: one for *each* document token t_i . MacAvaney et al. [44] explore several approaches to combine these contextualized vectors.

First, they extend the vanilla BERT ranker by making prediction on the average [CLS] token: $\frac{1}{m} \sum_{i=1}^m cls_i$. Second, they use contextualized embeddings as a direct replacement of context-free embeddings in the following neural architectures: KNRM [70], PACRR [33], and DRMM [23]. Third, they introduced a CEDR architecture where the [CLS] embedding is *additionally* incorporated into KNRM, PACCR, and DRMM in a model-specific way, which further boosts performance.

Non-parametric Model 1. Let $P(D|Q)$ denote a probability that a document D is relevant to the query Q . Using the Bayes rule, $P(D|Q)$ is convenient to re-write as $P(D|Q) \propto P(Q|D)P(D)$. Assuming a uniform prior for the document occurrence probability $p(D)$, one concludes that the relevance probability is proportional to $P(Q|D)$. Berger and Lafferty proposed to estimate this probability with a *term-independent* and *context-free* model known as Model 1 [4].

Let $T(q|d)$ be a probability that a query token q is a translation of a document token d and $P(d|D)$ is a probability that a token d is “generated” by a document D . Then, a probability that query Q is a translation of document D can be computed as a product of individual query term likelihoods as follows:

$$\begin{aligned} P(Q|D) &= \prod_{q \in Q} P(q|D) \\ P(q|D) &= \sum_{d \in D} T(q|d)P(d|D) \end{aligned} \quad (2)$$

The summation in Eq. 3 is over *unique* document tokens. The in-document term probability $P(d|D)$ is a maximum-likelihood estimate. Making the non-parametric Model 1 effective requires quite a few tricks. First, $P(q|D)$ —a likelihood of a query term q —is linearly combined with the collection probability $P(q|C)$ using a parameter λ [65, 71].¹

$$P(q|D) = (1 - \lambda) \left[\sum_{d \in D} T(q|d)P(d|D) \right] + \lambda P(q|C). \quad (3)$$

We take several additional measures to improve Model 1 effectiveness:

- We propose to create a parallel corpus by splitting documents and passages into small contiguous chunks whose length is comparable to query lengths;
- $T(q|d)$ are learned from a symmetrized corpus as proposed by Jeon et al. [35];
- We discard all translation probabilities $T(q|d)$ below an empirically found threshold of about 10^{-3} and keep at most 10^6 most frequent tokens;
- We make self-translation probabilities $T(t|t)$ to be equal to an empirically found positive value and rescale $T(t'|t)$ so that $\sum_{t'} T(t'|t) = 1$ as in [35, 65];

Our Neural Model 1. Let us rewrite Eq. 2 so that the inner summation is carried out over all document tokens rather than over the set of unique ones. This is particularly relevant for contextualized embeddings where embeddings of identical

¹ $P(q|C)$ is a maximum-likelihood estimate. For an out-of-vocabulary term q , $P(q|C)$ is set to a small number (e.g., 10^{-9}).

tokens are not guaranteed to be the same (and typically they are not):

$$P(Q|D) = \prod_{q \in Q} \sum_{i=1}^{|D|} \frac{T(q|d_i)}{|D|}. \quad (4)$$

We further propose to compute $T(q|d)$ in Eq. 4 by a simple and efficient neural network. Networks “consume” context-free or contextualized embeddings of tokens q and d and produces a value in the range $[0, 1]$. To incorporate a self translation probability—crucial for good convergence of the context-free model—we set $T(t|t) = p_{self}$ and multiply all other probabilities by $1 - p_{self}$. However, it was not practical to scale conditional probabilities to ensure that $\forall t_2 \sum_{t_1} T(t_1|t_2) = 1$. Thus, $T(t_1|t_2)$ is a similarity function, but not a true probability distribution. Note that—unlike CEDR [43]—we do not use the embedding of the [CLS] token.

We explored several approaches to neural parametrization of $T(t_1|t_2)$. Let $\text{embed}_q(t_1)$ and $\text{embed}_d(t_2)$ denote embeddings of query and document tokens, respectively. One of the simplest approaches is to learn separate embedding networks for queries and documents and use the scaled cosine similarity:

$$T(t_1|t_2) = 0.5\{\cos(\text{embed}_q(t_1), \text{embed}_d(t_2)) + 1\}.$$

However, this neural network is not sufficiently expressive and the resulting context-free Model 1 is inferior to the non-parametric Model 1 learned via EM. We then found that a key performance ingredient was a concatenation of embeddings with their Hadamard product, which we think helps the following layers discover better interaction features. We pass this combination through one or more fully-connected linear layer with RELUs [25] followed by a sigmoid:

$$\begin{aligned} T(q|d) &= \sigma(F_3(\text{relu}(F_2(\text{relu}(F_1([x_q, x_d, x_q \circ x_d])))))) \\ x_q &= P_q(\text{tanh}(\text{layer-norm}(\text{embed}_q(q)))) \\ x_d &= P_d(\text{tanh}(\text{layer-norm}(\text{embed}_d(d)))) \end{aligned}$$

where P_q , P_d , and F_i are fully-connected linear layers; $[x, y]$ is vector concatenation; layer-norm is layer normalization [2]; $x \circ y$ is the Hadamard product.

Neural Model 1 Sparsification/Export to Non-Parametric Format. We can pre-compute $T(t_1|t_2)$ for all pairs of vocabulary tokens, discard small values (below a threshold), and store the result as a sparse matrix. This format permits an extremely efficient execution on CPU (see results in Sect. 4.2).

4 Experiments

4.1 Setup

Data Sets. We experiment with MS MARCO collections, which include data for passage and document retrieval tasks [12, 49]. Each MS MARCO collection has

a large number of real user queries (see Table 1). To our knowledge, there are no other collections comparable to MS MARCO in this respect. The large set of queries is sampled from the log file of the search engine Bing. In that, data set creators ensured that all queries can be answered using a short text snippet. These queries are only sparsely judged (about one relevant passage per query). Sparse judgments are binary: Relevant documents have grade one and all other documents have grade zero.

In addition to large query sets with sparse judgments, we use two evaluation sets from TREC 2019/2020 deep learning tracks [12]. These query sets are quite small, but they have been thoroughly judged by NIST assessors *separately* for a document and a passage retrieval task. TREC NIST judgements range from zero (not-relevant) to three (perfectly relevant).

We randomly split publicly available training and validation sets into the following subsets: a small training set to train a linear fusion model (**train/fusion**), a large set to train neural models and non-parametric Model 1 (**train/modeling**), a development set (**development**), and a test set (MS MARCO **test**) containing at most 3K queries. Detailed data set statistics is summarized in Table 1. Note that the training subsets were obtained from the original training set, whereas the new development and test sets were obtained from the original development set. The leaderboard validation set is not publicly available.

We processed collections using Spacy 2.2.3 [30] to extract tokens (text words) and lemmas (lexemes) from text. The frequently occurring words and lemmas were filtered out using Indri’s list of stopwords [64], which was expanded to include a few contractions such as “n’t” and “ll”. Lemmas were indexed using Lucene 7.6. We also generated sub-word tokens, namely BERT word pieces [17, 69], using a HuggingFace Transformers library (version 0.6.2) [68]. We did *not* apply the stopwords list to BERT word pieces.

Basic Setup. We experimented on a Linux server equipped with a six-core (12 threads) i7-6800K 3.4 Ghz CPU, 125 GB of memory, and four GeForce GTX 1080 TI GPUs. We used the text retrieval framework FlexNeuART [8], which is implemented in Java. It employs Lucene 7.6 with a BM25 scorer [58] to generate an initial list of candidates, which can be further re-ranked using either traditional or neural re-rankers. The traditional re-rankers, including the non-parametric Model 1, are implemented in Java as well. They run in a *multi-threaded mode* (12 threads) and *fully* utilize the CPU. The neural rankers are implemented using PyTorch 1.4 [54] and Apache Thrift.² A neural ranker operates as a standalone *single-threaded* server. Our software is available online [8].³

Table 1. MS MARCO data set details

	Documents	Passages
# of documents	3.2M	8.8M
Avg. # of doc. lemmas	476.7	30.6
Avg. # of query lemmas	3.2	3.5
	# of queries	
Train/fusion	10K	20K
Train/modeling	357K	788.7K
Development	2500	20K
Test	2693	3000
TREC 2019	100	100
TREC 2020	100	100

² <https://thrift.apache.org/>.

³ <https://github.com/oaqa/FlexNeuART>.

Ranking speed is measured as the overall CPU/GPU *throughput*—rather than latency—per one *thousand* of documents/passages. Ranking accuracy is measured using the standard utility `trec_eval` provided by TREC organizers.⁴ Statistical significance is computed using a two-sided t-test with threshold 0.05.

All ranking models are applied to the candidate list generated by a tuned BM25 scorer [58]. BERT-based models re-rank 100 entries with highest BM25 scores: using a larger pool of candidates hurts both efficiency and accuracy. All other models, including the neural context-free Model 1 re-rank 1000 entries: Further increasing the number of candidates does not improve accuracy.

Training Models. Neural models are trained using a pairwise margin loss.⁵ Training pairs are obtained by combining known relevant documents with 20 negative examples selected from a set of top-500 candidates returned by Lucene. In each epoch, we randomly sample one positive and one negative example per query. BERT-based models first undergo a target-corpus pretraining [31] using a masked language modeling and next-sentence prediction objective [17]. Then, we train them for one epoch in a ranking task. We use batch size 16 simulated via gradient accumulation. Context-free Model 1 is trained from scratch for 32 epochs using batch size 32. The non-parametric Model 1 is trained for five epochs with MGIZA [53].⁶ Further increasing the number of epochs does not substantially improve results. MGIZA computes probabilities of spurious insertions (i.e., a translation from an empty word), but we discard them as in prior work [65].

We use a small weight decay (10^{-7}) and a warm-up schedule where the learning rate grows linearly from zero for 10–20% of the steps until it reaches the base learning rate [48, 63]. The optimizer is AdamW [42]. For BERT-based models we use different base rates for the fully-connected prediction head ($2 \cdot 10^{-4}$) and for the main Transformer layers ($2 \cdot 10^{-5}$). For the context-free Model 1 the base rate is $3 \cdot 10^{-3}$, which is decayed by 0.9 after each epoch. The learning rate is the same for all parameters.

The trained *neural* Model 1 is “exported” to a non-parametric format by precomputing all pairwise translation probabilities and discarding probabilities smaller than 10^{-4} . This sparsification/export procedure takes three minutes and the exported model is executed using the same Java code as the non-parametric Model 1. Each neural model and the sparsified Model 1 is trained and evaluated for five seeds. To this end, we compute the value for each query and seed and average query-specific values (over five seeds). All hyper-parameters are tuned on a development set.

Because context-free Model 1 rankers are not strong on their own, we evaluate them in a *fusion* mode. First, Model 1 is trained on `train/modeling`. Then we linearly combine a model score with the BM25 score [58]. Optimal weights are computed on a `train/fusion` subset using the coordinate ascent algorithm [45] from RankLib.⁷ To improve effectiveness of this linear fusion, we use Model 1

⁴ https://github.com/usnistgov/trec_eval.

⁵ We use the loss reduction type `sum`.

⁶ <https://github.com/moses-smt/mgiza/>.

⁷ <https://sourceforge.net/p/lemur/wiki/RankLib/>.

Table 2. Evaluation results: **bwps** denotes BERT word pieces, **lemm** denotes text lemmas, and **word** denotes original words. **NN-Model1** and **NN-Model1-exp** are the context-free neural Model 1 models: They use only **bwps**. **NN-Model1** runs on GPU whereas **NN-Model1-exp** runs on CPU. Ranking speed is throughput and not latency! Statistical significance is denoted by * and #. Hypotheses are explained in the main text.

	Documents				Passages			
	MS MARCO test	TREC 2019	TREC 2020	Rank. speed	MS MARCO test	TREC 2019	TREC 2020	Rank. speed
	MRR	NDCG@10		per 1K	MRR	NDCG@10		per 1K
Baselines								
BM25 (lemm)	0.270	0.544	0.524	0.8 ms	0.256	0.522	0.516	0.5 ms
BM25 (lemm)+BM25 (word)	0.274	0.544	0.523	2.5 ms	0.265	0.517	0.521	0.7 ms
BM25 (lemm)+BM25 (bwps)	0.283	0.528	0.537	2.2 ms	0.270	0.518	0.525	0.9 ms
BERT-vanilla (short)	0.387	0.655	0.623	39 s	0.426	0.686	0.684	15 s
BERT-vanilla (full)	0.376#	0.667	0.631	82 s				
BERT-CEDR-KRNM	0.387	0.665	0.649*	88 s	0.421*	0.682	0.675	16 ms
BERT-CEDR-DRMM	0.377*	0.667	0.636	120 s	0.425	0.688	0.685	30 s
BERT-CEDR-PACRR	0.392	0.670	0.652*	81 s	0.425	0.690	0.684	16 s
Our methods								
BM25 (lemm)+Model1 (word)	0.283*	0.548	0.535	13 ms	0.274*	0.522	0.567*	1.2 ms
BM25 (lemm)+Model1 (bwps)	0.284	0.557	0.525	33 ms	0.271	0.517	0.509	2.7 ms
BM25 (lemm)+NN-Model1-exp	0.307*	0.568	0.545	16 ms	0.298*	0.541*	0.581*	2.4 ms
BM25 (lemm)+NN-Model1	0.311*	0.566	0.541	3 s	0.300*	0.549*	0.587*	0.32 s
BERT-Model1 (short)	0.384	0.657	0.631	36 s	0.426	0.685	0.682	16 s
BERT-Model1 (full)	0.391#	0.666	0.637*	80 s				

log-scores *normalized* by the number of query words. In turn, BM25 scores are normalized by the sum of query-term IDF values (see [58] for the description of BM25 and IDF). As one of the baselines, we use a fusion of BM25 scores for different tokenization approaches (basically a multi-field BM25). Fusion weights are obtained via RankLib on **train/fusion**.

4.2 Results

Model Overview. We compare several models (see Table 2). First, we use BM25 scores [58] computed for the lemmatized text, henceforth, **BM25 (lemm)**. Second, we evaluate several variants of the context-free Model 1. The non-parametric Model 1 was trained for both original words and BERT word pieces: Respective models are denoted as **Model1 (word)** and **Model1 (bwps)**. The neural context-free Model 1—denoted as **NN-Model1**—was used only with BERT word pieces. This model was sparsified and exported to a non-parametric format (see Sect. 3), which runs efficiently on a CPU. We denote it as **NN-Model1-exp**. Note that context-free Model 1 rankers are not strong on their own, thus, we evaluate them in a *fusion* mode by combining their scores with **BM25 (lemm)**.

Crucially, all context-free models incorporate exact term-matching signal via either the self-translation probability or via explicit smoothing with a word collection probability (see Eq. 3). Thus, these models should be compared not only with BM25, but also with the fusion model incorporating BM25 scores for original words or BERT word pieces. We denote these baselines as **BM25 (lemm)+BM25 (word)** and **BM25 (lemm)+BM25 (bwps)**, respectively.

As we describe in Sect. 3, our contextualized Model 1 applies the neural Model 1 layer to the contextualized embeddings produced by BERT. We denote

this model as `BERT-Model1`. Due to the limitation of existing pretrained Transformer models, long documents need to be split into chunks each of which is processed, i.e., contextualized, separately. This is done in `BERT-Model1 (full)`, `BERT-vanilla (full)`, and `BERT-CEDR [44]` models. These models operate on (mostly) complete documents: For efficiency reasons we nevertheless use only the first 1431 tokens (three BERT chunks). Another approach is to make predictions on much shorter (one BERT chunk) fragments [15]. This is done in `BERT-Model1 (short)` and `BERT-vanilla (short)`. In the passage retrieval task, all passages are short and no truncation or chunking is needed. Note that we use a *base*, i.e., a 12-layer Transformer [67] model, since it is more practical than a 24-layer `BERT-large` and performs at par with `BERT-large` on MS MARCO data [29].

We tested several hypotheses using a two-sided t-test:

- `BM25 (1emm)+ Model1 (word)` is the same as `BM25 (1emm)+ BM25 (word)`;
- `BM25 (1emm)+ Model1 (bwps)` is the same as `BM25 (1emm)+ BM25 (bwps)`;
- `BERT-Model1 (full)` is the same as `BERT-vanilla (short)`;
- For each `BERT-CEDR` model, we test if it is the same as `BERT-vanilla (short)`;
- `BERT-vanilla (full)` is the same as `BERT-vanilla (short)`;
- `BERT-Model1 (full)` is the same as `BERT-Model1 (short)`;

The main purpose of these tests is to assess if special aggregation layers (including the neural Model 1) can be more accurate compared to models that run on truncated documents. In Table 2 statistical significance is indicated by a special symbol: the last two hypotheses use #; all other hypotheses use *.

Discussion of Results. The results are summarized in Table 2. First note that there is less consistency in results on TREC 2019/2020 sets compared to MS MARCO test sets. In that, some statistically significant differences (on MS MARCO test) “disappear” on TREC 2019/2020. TREC 2019/2020 query sets are quite small and its more likely (compared to MS MARCO test) to obtain spurious results. Furthermore, the fusion model `BM25 (1emm)+ Model1 (bwps)` is either worse than the baseline model `BM25 (1emm)+ BM25 (bwps)` or the difference is not significant. `BM25 (1emm)+ Model1 (word)` is mostly better than the respective baseline, but the gain is quite small. In contrast, the fusion of the neural Model 1 with BM25 scores for BERT word pieces is more accurate on all the query sets. On the MS MARCO test sets it is 15–17% better than `BM25 (1emm)`. These differences are significant on both MS MARCO test sets as well as on TREC 2019/2020 tests sets for the passage retrieval task. Sparsification of the neural Model 1 leads only to a small (0.6–1.3%) loss in accuracy. In that, the sparsified model—executed on a CPU—is more than 10^3 times faster than BERT-based rankers, which run on a GPU. It is $5 \times 10^3 \times$ faster in the case of passage retrieval. In contrast, on a GPU, the fastest neural model KNNRM is only 500 times faster than vanilla BERT [28] (also for passage retrieval). For large candidate sets computation of Model 1 scores can be further sped up (Sect. 3.1.2.1 [6]). Thus, `BM25 (1emm)+NN-Model1-exp` can be useful at the candidate generation stage.

We also compared BERT-based neural Model 1 with `BERT-CEDR` and `BERT-vanilla` models on the MS MARCO test set for the document retrieval

task. By comparing BERT-vanilla (short), BERT-Model1 (short), and BERT-Model1 (full) we can see that the neural Model 1 layer entails virtually no efficiency or accuracy loss. In fact, BERT-Model1 (full) is 1.8% and 1% better than BERT-Model1 (short) and BERT-vanilla (short), respectively. Yet, only the former difference is statistically significant.

Furthermore, the same holds for BERT-CEDR-PACRR, which was shown to outperform BERT-vanilla by MacAvaney et al. [44]. In our experiments it is 1% better than BERT-vanilla (short), but the difference is neither substantial nor statistically significant. This does not invalidate results of MacAvaney et al. [44]: They compared BERT-CEDR-PACRR only with BERT-vanilla (full), which makes predictions on the averaged [CLS] embeddings. However, in our experiments, this model is noticeably worse (by 4.2%) than BERT-vanilla (short) and the difference is statistically significant. We think that obtaining more conclusive evidence about the effectiveness of aggregation layers requires a different data set where relevance is harder to predict from a truncated document.

Leaderboard Submissions. We combined BERT-Model1 with the strong first-stage pipeline, which uses Lucene to index documents expanded with doc2query [51, 52] and re-ranks them using a mix of traditional and NN-Model1-exp scores (our exported neural Model 1). This first-stage pipeline is about as effective as the Conformer-Kernel model [47]. The combination model achieved the top place on a well-known leaderboard in November and December 2020. Furthermore, using the non-parametric Model 1, we produced the best traditional run in December 2020, which outperformed several neural baselines [7].

5 Conclusion

We study a neural Model 1 combined with a context-free or contextualized embedding network and show that such a combination has benefits to efficiency, effectiveness, and interpretability. To our knowledge, the context-free neural Model 1 is the only neural model that can be sparsified to run efficiently on a CPU (up to 5×10^3 × faster than BERT on a GPU) without expensive index-time precomputation or query-time operations on large tensors. We hope that effectiveness of this approach can be further improved, e.g., by designing a better parametrization of conditional translation probabilities.

References

1. MS MARCO leaderboard. <https://microsoft.github.io/msmarco/>
2. Ba, L.J., Kiros, J.R., Hinton, G.E.: Layer normalization, CoRR abs/1607.06450 (2016)
3. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations, ICLR 2015 (2015)
4. Berger, A., Lafferty, J.: Information retrieval as statistical translation. In: Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 222–229 (1999)

5. Bourlard, H., Bourlard, H.A., Morgan, N.: Connectionist Speech Recognition. A Hybrid Approach, vol. 247. Springer, New York (1994). <https://doi.org/10.1007/978-1-4615-3210-1>
6. Boytsov, L.: Efficient and Accurate Non-Metric k-NN Search with Applications to Text Matching. Ph.D. thesis, Carnegie Mellon University (2018)
7. Boytsov, L.: Traditional IR rivals neural models on the MS MARCO document ranking leaderboard (2020)
8. Boytsov, L., Nyberg, E.: Flexible retrieval with NMSLIB and FlexNeuART. In: Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS), pp. 32–43 (2020)
9. Brown, P.F., Pietra, S.D., Pietra, V.J.D., Mercer, R.L.: The mathematics of statistical machine translation: parameter estimation. *Comput. Linguist.* **19**(2), 263–311 (1993)
10. Büttcher, S., Clarke, C.L., Cormack, G.V.: Information Retrieval: Implementing and Evaluating Search Engines. MIT Press, Cambridge (2016)
11. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., Kuksa, P.: Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **12**, 2493–2537 (2011)
12. Craswell, N., Mitra, B., Yilmaz, E., Campos, D., Voorhees, E.M.: Overview of the TREC 2019 deep learning track, CoRR abs/2003.07820 (2020)
13. Croft, W.B., Metzler, D., Strohman, T.: Search Engines: Information Retrieval in Practice. Addison-Wesley, Reading (2010)
14. Dai, Z., Callan, J.: Context-aware sentence/passage term importance estimation for first stage retrieval, CoRR abs/1910.10687 (2019)
15. Dai, Z., Callan, J.: Deeper text understanding for IR with contextual neural language modeling. In: SIGIR, pp. 985–988. ACM (2019)
16. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. Ser. B (Methodol.)* **39**(1), 1–22 (1977)
17. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding, pp. 4171–4186 (2019)
18. Furnas, G.W., Landauer, T.K., Gomez, L.M., Dumais, S.T.: The vocabulary problem in human-system communication. *Commun. ACM* **30**(11), 964–971 (1987)
19. Ganguly, D., Roy, D., Mitra, M., Jones, G.J.: Word embedding based generalized language model for information retrieval. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 795–798 (2015)
20. Gao, L., Dai, Z., Callan, J.: EARL: speedup transformer-based rankers with pre-computed representation, CoRR abs/2004.13313 (2020)
21. Gers, F.A., Schmidhuber, J., Cummins, F.A.: Learning to forget: continual prediction with LSTM. *Neural Comput.* **12**(10), 2451–2471 (2000)
22. Goldberg, Y.: A primer on neural network models for natural language processing. *J. Artif. Intell. Res.* **57**, 345–420 (2016)
23. Guo, J., Fan, Y., Ai, Q., Croft, W.B.: A deep relevance matching model for ad-hoc retrieval. In: CIKM, pp. 55–64. ACM (2016)
24. Guo, J., et al.: A deep look into neural ranking models for information retrieval. *Inf. Process. Manage.* **57**, 102067 (2019)
25. Hahnloser, R.H.R.: On the piecewise analysis of networks of linear threshold neurons. *Neural Netw.* **11**(4), 691–697 (1998)

26. Hinton, G., et al.: Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Sig. Process. Mag.* **29**(6), 82–97 (2012)
27. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
28. Hofstätter, S., Hanbury, A.: Let’s measure run time! extending the IR replicability infrastructure to include performance aspects. In: *OSIRRC@SIGIR*. *CEUR Workshop Proceedings*, vol. 2409, pp. 12–16. *CEUR-WS.org* (2019)
29. Hofstätter, S., Zlabinger, M., Hanbury, A.: Interpretable & time-budget-constrained contextualization for re-ranking. In: *Frontiers in Artificial Intelligence and Applications (ECAI)*, vol. 325, pp. 513–520. *IOS Press* (2020)
30. Honnibal, M., Montani, I.: *spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing* (2017, to appear)
31. Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. In: *ACL*, vol. 1, pp. 328–339. *Association for Computational Linguistics* (2018)
32. Huang, Z., Xu, W., Yu, K.: Bidirectional LSTM-CRF models for sequence tagging, *CoRR abs/1508.01991* (2015)
33. Hui, K., Yates, A., Berberich, K., de Melo, G.: Co-PACRR: a context-aware neural IR model for ad-hoc retrieval. In: *WSDM*, pp. 279–287. *ACM* (2018)
34. Jain, S., Wallace, B.C.: Attention is not explanation. In: *NAACL-HLT*, vol. 1, pp. 3543–3556. *Association for Computational Linguistics* (2019)
35. Jeon, J., Croft, W.B., Lee, J.H.: Finding similar questions in large question and answer archives. In: *CIKM*, pp. 84–90. *ACM* (2005)
36. Khattab, O., Potts, C., Zaharia, M.: Relevance-guided supervision for OpenQA with ColBERT, *CoRR abs/2007.00814* (2020)
37. Khattab, O., Zaharia, M.: ColBERT: efficient and effective passage search via contextualized late interaction over BERT. In: *SIGIR*, pp. 39–48. *ACM* (2020)
38. Lavrenko, V., Choquette, M., Croft, W.B.: Cross-lingual relevance models. In: *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 175–182 (2002)
39. Li, C., Yates, A., MacAvaney, S., He, B., Sun, Y.: PARADE: passage representation aggregation for document reranking, *CoRR abs/2008.09093* (2020)
40. Lin, J.: The neural hype and comparisons against weak baselines. In: *ACM SIGIR Forum*, vol. 52, pp. 40–51. *ACM*, New York (2019)
41. Lipton, Z.C.: The mythos of model interpretability. *Commun. ACM* **61**(10), 36–43 (2018)
42. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint [arXiv:1711.05101](https://arxiv.org/abs/1711.05101)* (2017)
43. MacAvaney, S., Nardini, F.M., Perego, R., Tonello, N., Goharian, N., Frieder, O.: Expansion via prediction of importance with contextualization. In: *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1573–1576. *ACM* (2020)
44. MacAvaney, S., Yates, A., Cohan, A., Goharian, N.: CEDR: contextualized embeddings for document ranking. In: *SIGIR*, pp. 1101–1104. *ACM* (2019)
45. Metzler, D., Croft, W.B.: Linear feature-based models for information retrieval. *Inf. Retr.* **10**(3), 257–274 (2007). <https://doi.org/10.1007/s10791-006-9019-z>
46. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *NIPS*, pp. 3111–3119 (2013)

47. Mitra, B., Hofstätter, S., Zamani, H., Craswell, N.: Conformer-kernel with query term independence for document retrieval, CoRR abs/2007.10434 (2020)
48. Mosbach, M., Andriushchenko, M., Klakow, D.: On the stability of fine-tuning BERT: misconceptions, explanations, and strong baselines, CoRR abs/2006.04884 (2020)
49. Nguyen, T., et al.: MS MARCO: a human generated MACHine Reading COMprehension dataset (November 2016)
50. Nogueira, R., Cho, K.: Passage re-ranking with BERT, CoRR abs/1901.04085 (2019)
51. Nogueira, R., Lin, J.: From doc2query to docTTTTTquery. An MS MARCO passage retrieval task [1] μ publication (2019)
52. Nogueira, R., Yang, W., Lin, J., Cho, K.: Document expansion by query prediction, CoRR abs/1904.08375 (2019)
53. Och, F.J., Ney, H.: A systematic comparison of various statistical alignment models. *Comput. Linguist.* **29**(1), 19–51 (2003)
54. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems*, pp. 8026–8037 (2019)
55. Peters, M.E., et al.: Deep contextualized word representations. In: *Proceedings of NAACL-HLT*, pp. 2227–2237 (2018)
56. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding with unsupervised learning. Technical report, OpenAI (2018)
57. Riezler, S., Vasserman, A., Tsochantaridis, I., Mittal, V.O., Liu, Y.: Statistical machine translation for query expansion in answer retrieval. In: *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, ACL 2007* (2007)
58. Robertson, S.: Understanding inverse document frequency: on theoretical arguments for IDF. *J. Doc.* **60**(5), 503–520 (2004)
59. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **1**(5), 206–215 (2019)
60. Rush, A.M.: The annotated transformer. In: *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pp. 52–60 (2018)
61. Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. *IEEE Trans. Sig. Process.* **45**(11), 2673–2681 (1997)
62. Serrano, S., Smith, N.A.: Is attention interpretable? In: *ACL*, vol. 1, pp. 2931–2951. Association for Computational Linguistics (2019)
63. Smith, L.N.: Cyclical learning rates for training neural networks. In: *WACV*, pp. 464–472. IEEE Computer Society (2017)
64. Strohman, T., Metzler, D., Turtle, H., Croft, W.B.: Indri: a language-model based search engine for complex queries (2005). <http://ciir.cs.umass.edu/pubfiles/ir-407.pdf>. Accessed April 2017
65. Surdeanu, M., Ciaramita, M., Zaragoza, H.: Learning to rank answers to non-factoid questions from web collections. *Comput. Linguist.* **37**(2), 351–383 (2011)
66. Tay, Y., Dehghani, M., Bahri, D., Metzler, D.: Efficient transformers: a survey, CoRR abs/2009.06732 (2020)
67. Vaswani, A., et al.: Attention is all you need. In: *NIPS*, pp. 5998–6008 (2017)
68. Wolf, T., et al.: Huggingface’s transformers: state-of-the-art natural language processing. arXiv abs/1910.03771 (2019)
69. Wu, Y., et al.: Google’s neural machine translation system: bridging the gap between human and machine translation. CoRR abs/1609.08144 (2016)

70. Xiong, C., Dai, Z., Callan, J., Liu, Z., Power, R.: End-to-end neural ad-hoc ranking with kernel pooling. In: SIGIR, pp. 55–64. ACM (2017)
71. Xue, X., Jeon, J., Croft, W.B.: Retrieval models for question and answer archives. In: SIGIR, pp. 475–482 (2008)
72. Yilmaz, Z.A., Wang, S., Yang, W., Zhang, H., Lin, J.: Applying BERT to document retrieval with birch. In: EMNLP/IJCNLP, vol. 3, pp. 19–24. Association for Computational Linguistics (2019)
73. Zbib, R., et al.: Neural-network lexical translation for cross-lingual IR from text and speech. In: SIGIR, pp. 645–654. ACM (2019)
74. Zhao, L., Callan, J.: Term necessity prediction. In: Huang, J., Koudas, N., Jones, G.J.F., Wu, X., Collins-Thompson, K., An, A. (eds.) Proceedings of the 19th ACM Conference on Information and Knowledge Management, CIKM 2010, pp. 259–268. ACM (2010)
75. Zuccon, G., Koopman, B., Bruza, P., Azzopardi, L.: Integrating and evaluating neural word embeddings in information retrieval. In: Proceedings of the 20th Australasian Document Computing Symposium, pp. 1–8 (2015)