
CONFORMER-KERNEL WITH QUERY TERM INDEPENDENCE FOR DOCUMENT RETRIEVAL

Bhaskar Mitra
Microsoft, UCL
bmitra@microsoft.com

Sebastian Hofstätter
TU Wien
s.hofstatter@tuwien.ac.at

Hamed Zamani and Nick Craswell
Microsoft
{hazamani, nickcr}@microsoft.com

ABSTRACT

The Transformer-Kernel (TK) model has demonstrated strong reranking performance on the TREC Deep Learning benchmark—and can be considered to be an efficient (but slightly less effective) alternative to BERT-based ranking models. In this work, we extend the TK architecture to the full retrieval setting by incorporating the query term independence assumption. Furthermore, to reduce the memory complexity of the Transformer layers with respect to the input sequence length, we propose a new Conformer layer. We show that the Conformer’s GPU memory requirement scales linearly with input sequence length, making it a more viable option when ranking long documents. Finally, we demonstrate that incorporating explicit term matching signal into the model can be particularly useful in the full retrieval setting. We present preliminary results from our work in this paper.

Keywords Deep learning · Neural information retrieval · Ad-hoc retrieval

1 Introduction

In the inaugural year of the TREC Deep Learning track [Craswell et al., 2019], Transformer-based [Vaswani et al., 2017] ranking models demonstrated substantial improvements over traditional information retrieval (IR) methods. Several of these approaches—*e.g.*, [Yilmaz et al., 2019, Yan et al., 2019]—employ BERT [Devlin et al., 2018], with large-scale pretraining, as their core architecture. Diverging from the trend of BERT-scale models, Hofstätter et al. [2020b] propose the Transformer-Kernel (TK) model with few key distinctions: (i) TK uses a shallower model with only two Transformer layers, (ii) The parameters of the model are randomly initialized prior to training (skipping the computation-intensive pretraining step), and finally (iii) TK encodes the query and the document independently of each other allowing for offline precomputations for faster response times. Consequently, TK achieves competitive performance at a fraction of the training and inference cost of its BERT-based peers.

Notwithstanding these efficiency gains, the TK model shares two critical drawbacks with other Transformer-based models. Firstly, the memory complexity of the self-attention layers is quadratic $\mathcal{O}(n^2)$ with respect to the length n of the input sequence. This restricts the number of document terms that the model can inspect under fixed GPU memory budget. A trivial workaround involves inspecting only the first k terms of the document. This approach can not only negatively impact retrieval quality, but has been shown to specifically under-retrieve longer documents [Hofstätter et al., 2020a]. Secondly, in any real IR system, it is impractical to evaluate every document in the collection for every query—and therefore these systems typically either enforce some sparsity property to drastically narrow down the set of documents that should be evaluated or find ways to prioritize the candidates for evaluation. TK employs a nonlinear matching function over query-document pairs. Therefore, it is not obvious how the TK function can be directly used to retrieve from the full collection without exhaustively comparing every document to the query. This restricts TK’s scope of application to late stage reranking of smaller candidate sets that may have been identified by simpler retrieval models.

In this work, we extend the TK architecture to enable direct retrieval from the full collection of documents. Towards that goal, we incorporate three specific changes:

1. To scale to long document text, we replace each instance of the Transformer layer with a novel Conformer layer whose memory complexity is $\mathcal{O}(n \times d_{\text{key}})$, instead of $\mathcal{O}(n^2)$.

2. To enable fast retrieval with TK, we incorporate the query term independence assumption [Mitra et al., 2019] into the architecture.
3. And finally, as Mitra et al. [2016, 2017] point out, lexical term matching can complement latent matching models, and the combination can be particularly effective when retrieving from the full collection of candidates. So, we extend TK with an explicit term matching submodel to minimize the impact of false positive matches in the latent space.

We describe the full model and present preliminary results from our work in this paper.

2 Related work

2.1 Scaling self-attention to long text

The self-attention layer, as proposed by Vaswani et al. [2017], can be described as follows:

$$\text{Self-Attention}(Q, K, V) = \Phi\left(\frac{QK^\top}{\sqrt{d_k}}\right) \cdot V \quad (1)$$

Where, $Q \in \mathbb{R}^{n \times d_{\text{key}}}$, $K \in \mathbb{R}^{n \times d_{\text{key}}}$, and $V \in \mathbb{R}^{n \times d_{\text{value}}}$ are the query, key, and value matrices—and d_{key} and d_{value} are the dimensions of the key and value embeddings, respectively, and n is the length of the input sequence. We use Φ to denote the softmax operation applied along the last dimension of the input matrix.

The quadratic $\mathcal{O}(n^2)$ memory complexity of self-attention is a direct consequence of the component $QK^\top$ that produces a matrix of size $n \times n$. Recently, an increasing number of different approaches have been proposed in the literature to get around this quadratic complexity. Broadly speaking, most of these approaches can be classified as either: (i) Restricting self-attention to smaller windows over the input sequence which results in a memory complexity of $\mathcal{O}(n \times m)$, where m is the window size—e.g., [Parmar et al., 2018, Dai et al., 2019, Yang et al., 2019, Sukhbaatar et al., 2019], or (ii) Operating under the assumption that the attention matrix is low rank r and hence finding alternatives to explicitly computing the $QK^\top$ matrix to achieve a complexity of $\mathcal{O}(n \times r)$ —e.g., [Kitaev et al., 2019, Roy et al., 2020, Tay et al., 2020, Wang et al., 2020], or (iii) A hybrid of both approaches—e.g., [Child et al., 2019, Beltagy et al., 2020, Wu et al., 2020]. In the IR literature, recently Hofstätter et al. [2020a] have extended the TK model to longer text using the local window-based attention approach. Other more general approaches to reducing the memory footprint of very deep models, such as model parallelization have also been extended to Transformer models [Shoeybi et al., 2019]. For more general primer on self-attention and Transformer architectures, we point the reader to Weng [2018, 2020].

2.2 Full retrieval with deep models

Efficient retrieval using complex machine learned relevance functions is an important challenge in neural IR [Mitra and Craswell, 2018, Guo et al., 2019]. One family of approaches involves the dual encoder architecture where the query and document are encoded independently of each other, and efficient retrieval is achieved using approximate nearest-neighbour search [Lee et al., 2019, Chang et al., 2020, Karpukhin et al., 2020, Ahmad et al., 2019, Khattab and Zaharia, 2020] or by employing other data structures, such as learning an inverted index based on latent representations [Zamani et al., 2018]. Precise matching of terms or concepts may be difficult using query-independent latent document representations [Luan et al., 2020], and therefore these models are often combined with explicit term matching methods [Nalisnick et al., 2016, Mitra et al., 2017]. Xiong et al. [2020] have recently demonstrated that the training data distribution can also significantly influence the performance of dual encoder models under the full retrieval setting. Auxilliary optimization objectives can also help guide the training of latent matching models to find solutions that emphasize more precise matching of terms and concepts [Rosset et al., 2019].

An alternative approach assumes query term independence (QTI) in the design of the neural ranking model [Mitra et al., 2019]. For these family of models, the estimated relevance score $S_{q,d}$ is factorized as a sum of the estimated relevance of the document to each individual query term.

$$S_{q,d} = \sum_{t \in q} s_{t,d} \quad (2)$$

Readers should note that the QTI assumption is already baked into several classical IR models, like BM25 [Robertson et al., 2009]. Relevance models with QTI assumption can be used to precompute all term-document scores offline. The precomputed scores can be subsequently leveraged for efficient search using inverted-index data structures.

Several recent neural IR models [Mitra et al., 2019, Dai and Callan, 2019b,a, Mackenzie et al., 2020, Dai and Callan, MacAvaney et al., 2020] that incorporate the QTI assumption have obtained promising results under the full retrieval setting. Document expansion based methods [Nogueira et al., 2019b,a], using large neural language models, can also be classified as part of this family of approaches, assuming the subsequent retrieval step employs a traditional QTI model like BM25. In all of these approaches, the focus of the machine learned function is to estimate the impact score of the document with respect to individual terms in the vocabulary, which can be precomputed offline during index creation.

An obvious alternative to document expansion based methods is to use the neural model to reformulate the query [Nogueira and Cho, 2017, Van Gysel et al., 2017, Ma et al., 2020]—although these approaches have not yet demonstrated retrieval performance that can be considered competitive to other methods considered here.

Finally, when the relevance of items are known, or a reliable proxy metric exists, machine learned policies [Kraska et al., 2018, Oosterhuis et al., 2018, Rosset et al., 2018] can also be effective for efficient search over indexes but these methods are not directly relevant to our current discussion.

3 Conformer-Kernel with QTI

We begin by briefly describing the original TK model as outlined in Fig 1a. The initial word embedding layer in TK maps both query and document to their respective sequences of term embeddings. These sequences are then passed through one or more stacked Transformer layers to derive contextualized vector representations for query and document terms. The learnable parameters of both query and document encoders are shared—which includes the initial term embeddings as well as the Transformer layers. Based on the contextualized term embeddings, TK creates an interaction matrix X , such that X_{ij} is the cosine similarity between the contextualized embeddings of the i^{th} query term q_i and the j^{th} document term d_j .

$$X_{ij} = \cos(\vec{v}_{q_i}, \vec{v}_{d_j}) \quad (3)$$

The Kernel-Pooling stage then creates k distinct features per query term as follows:

$$K_{ik} = \log \sum_j^{|d|} \exp\left(-\frac{(X_{ij} - \mu_k)^2}{2\sigma^2}\right) \quad (4)$$

Finally, the query-document relevance is estimated by a nonlinear function—typically implemented as stacked feedforward layers—over these features. Next, we describe the proposed changes to this base architecture.

3.1 Conformer

In Section 2.1, we note that the quadratic memory complexity of the self-attention layers *w.r.t.* the length of the input sequence is a direct result of explicitly computing the attention matrix $QK^\top \in \mathbb{R}^{n \times n}$. In this work, we propose a new separable self-attention layer that allows us to avoid instantiating the full term-term attention matrix as follows:


$$\text{Separable-Self-Attention}(Q, K, V) = \Phi(Q) \cdot A \quad (5)$$

$$\text{where, } A = \Phi(K^\top) \cdot V \quad (6)$$


As previously, Φ denotes the softmax operation along the last dimension of the input matrix. Note that, however, in this separable self-attention mechanism, the softmax operation is employed twice: (i) $\Phi(Q)$ computes the softmax along the d_{key} dimension, and (ii) $\Phi(K^\top)$ computes the softmax along the n dimension. By computing $A \in \mathbb{R}^{d_{\text{key}} \times d_{\text{value}}}$ first, we avoid explicitly computing the full term-term attention matrix. The memory complexity of the separable self-attention layer is $\mathcal{O}(n \times d_{\text{key}})$, which is a significant improvement when $d_{\text{key}} \ll n$.

We modify the standard Transformer block as follows:

1. We replace the standard self-attention layer with the more memory efficient separable self-attention layer.
2. Furthermore, we apply grouped convolution before the separable self-attention layers to better capture the local context based on the window of neighbouring terms.



(a) Transformer-Kernel (TK)



(b) Conformer-Kernel (CK) with QTI

Figure 1: A comparison of the TK and the proposed CK-with-QTI architectures. In addition to replacing the Transformer layers with Conformers, the latter also simplifies the query encoding to non-contextualized term embedding lookup and incorporates a windowed Kernel-Pooling based aggregation that is employed independently per query term.

We refer to this combination of grouped convolution and Transformer with separable self-attention as a Conformer. We incorporate the Conformer layer into TK as a direct replacement for the existing Transformer layers and name the new architecture as a Conformer-Kernel (CK) model. In relation to handling long input sequences, we also replace the standard Kernel-Pooling with windowed Kernel-Pooling [Hofstätter et al., 2020a] in our proposed architecture.

3.2 Query term independence assumption

To incorporate the QTI assumption into TK, we make a couple of simple modifications to the original architecture. Firstly, we simplify the query encoder by getting rid of all the Transformer layers and only considering the non-contextualized embeddings for the query terms. Secondly, instead of applying the aggregation function over the full interaction matrix, we apply it to each row of the matrix individually, which corresponds to individual query terms. The scalar outputs from the aggregation function are linearly combined to produce the final query-document score. These proposed changes are shown in Fig 1b.

3.3 Explicit term matching

We adopt the Duet [Nanni et al., 2017, Mitra and Craswell, 2019b,a] framework wherein the term-document score is a linear combination of outputs from a latent and an explicit matching models.

$$s_{t,d} = w_1 \cdot \text{BN}(s_{t,d}^{(\text{latent})}) + w_2 \cdot \text{BN}(s_{t,d}^{(\text{explicit})}) + b \quad (7)$$

Where, $\{w_1, w_2, b\}$ are learnable parameters and BN denotes the BatchNorm operation [Ioffe and Szegedy, 2015].

$$\text{BN}(x) = \frac{x - \mathbb{E}[x]}{\sqrt{\text{Var}[x]}} \quad (8)$$

We employ the CK model to compute $s_{t,d}^{(\text{latent})}$ and define a new lexical matching function modeled on BM25 for $s_{t,d}^{(\text{explicit})}$.

$$s_{t,d}^{(\text{explicit})} = \text{IDF}_t \cdot \frac{\text{BS}(\text{TF}_{t,d})}{\text{BS}(\text{TF}_{t,d}) + \text{ReLU}(w_{\text{dlen}} \cdot \text{BS}(|d|) + b_{\text{dlen}}) + \epsilon} \quad (9)$$

Where, IDF_t , $\text{TF}_{t,d}$, and $|d|$ denote the inverse-document frequency of the term t , the term-frequency of t in document d , and the length of the document, respectively. The w_{dlen} and b_{dlen} are the only two learnable parameters of this submodel and ϵ is a small constant added to prevent a divide-by-zero error. The BatchScale (BS) operation is defined as follows:

$$\text{BS}(x) = \frac{x}{\mathbb{E}[x] + \epsilon} \quad (10)$$

4 Experiments

4.1 Task and data

We conduct preliminary experiments on the document retrieval benchmark provided as part of the TREC Deep Learning track [Craswell et al., 2019]. The benchmark is based on the MS MARCO dataset [Bajaj et al., 2016] and provides a collection of 3, 213, 835 documents and a training dataset with 384, 597 positively labeled query-document pairs. Recently, the benchmark also made available a click log dataset, called ORCAS [Craswell et al., 2020], that can be employed as an additional document description field. We refer the reader to the track website¹ for further details about the benchmark.

Because we are interested in the full ranking setting, we do not make use of the provided document candidates and instead use the proposed model to retrieve from the full collection. We compare different runs based on following three metrics: mean reciprocal rank (MRR) [Craswell, 2009], normalized discounted cumulative gain (NDCG) [Järvelin and Kekäläinen, 2002], and normalized cumulative gain (NCG) [Rosset et al., 2018].

¹<https://microsoft.github.io/TREC-2020-Deep-Learning/>

Table 1: Full retrieval results based on the TREC 2019 Deep Learning track test set.

Model	MRR	NDCG@10	NCG@100
Non-neural baselines			
BM25+RM3 run with best NDCG@10	0.807	0.549	0.559
Non-neural run with best NDCG@10	0.872	0.561	0.560
Neural baselines			
DeepCT run with best NDCG@10	0.872	0.554	0.498
BERT-based document expansion + reranking run with best NCG@10	0.899	0.646	0.637
BERT-based document expansion + reranking run with best NDCG@10	0.961	0.726	0.580
Our models			
Conformer-Kernel	0.845	0.554	0.464
Conformer-Kernel + learned BM25	0.906	0.603	0.533
Conformer-Kernel + learned BM25 + ORCAS field	0.898	0.620	0.547

4.2 Model training

We consider the first 20 terms for every query and the first 4000 terms for every document. When incorporating the ORCAS data as an additional document field, we limit the maximum length of the field to 2000 terms. We pretrain the word embeddings using the word2vec [Mikolov et al., 2013a,b,c] implementation in FastText [Joulin et al., 2016]. We use a concatenation of the IN and OUT embeddings [Nalisnick et al., 2016, Mitra et al., 2016] from word2vec to initialize the embedding layer parameters. The document encoder uses 2 Conformer layers and we set all the hidden layer sizes to 256. We set the window size for the grouped convolution layers to 31 and the number of groups to 32. Correspondingly, we also set the number of attention heads to 32. We set the number of kernels k to 10. For windowed Kernel-Pooling, we set the window size to 300 and the stride to 100. Finally, we set the dropout rate to 0.2. For further details, please refer to the publicly released model implementation in PyTorch.² All models are trained on four Tesla P100 GPUs, with 16 GB memory each, using data parallelism.

We train the model using the RankNet objective [Burges et al., 2005]. For every positively labeled query-document pair in the training data, we randomly sample one negative document from the provided top 100 candidates corresponding to the query and two negative documents from the full collection. In addition to making pairs between the positively labeled document and the three negative documents, we also create pairs between the negative document sampled from the top 100 candidates and those sampled from the full collection, treating the former as more relevant. This can be interpreted as incorporating a form of weak supervision [Dehghani et al., 2017] as the top candidates were previously generated using a traditional IR function.

5 Results

Table 1 presents our main experiment results. As specified earlier, we evaluate our models on the full ranking setting without any explicit reranking step. The full model—with both Conformer-Kernel and explicit matching submodel—performs significantly better on NDCG@10 and MRR compared to the best traditional runs from the 2019 edition of the track. The model also outperforms the DeepCT baseline which is a QTI-based baseline using BERT. The other BERT-based baselines outperform our model by significant margins. We believe this observation should motivate future exploration on how to incorporate pretraining in the Conformer-Kernel model. Finally, we also notice improvements from incorporating the ORCAS data as an additional document descriptor field.

To demonstrate how the GPU memory consumption scales with respect to input sequence length, we plot the peak memory, across all four GPUs, for our proposed architecture using Transformer and Conformer layers, respectively, keeping all other hyperparameters and architecture choices fixed. Fig 2 shows the GPU memory requirement grows linearly with increasing sequence length for the Conformer, while quadratically when Transformer layers are employed.

6 Discussion and future work

The proposed CK-with-QTI architecture provides several advantages, with respect to inference cost, compared to its BERT-based peers. In addition to a shallower model and more memory-efficient Conformer layers, the model allows for offline pre-encoding of documents during indexing. It is notable, that the document encoder, containing the stacked Conformer layers, is the computationally costliest part of the model. In the proposed architecture, the document

²<https://github.com/bmitra-msft/TREC-Deep-Learning-Quick-Start>



Figure 2: Comparison of peak GPU Memory Usage in MB, across all four GPUs, when employing Transformers vs. Conformers in our proposed architecture. For the Transformer-based model, we only plot till sequence length of 512, because for longer sequences we run out of GPU memory when using Tesla P100s with 16 GB of memory.

encoder needs to be evaluated only once per every document in the collection. This is in contrast to once per every query-document pair in the case of BERT-based ranking models that accepts a concatenation of query and document as input [Nogueira and Cho, 2019], and once per every term-document pair in the case of BERT-based ranking models with QTI [Mitra et al., 2019].

While the present study demonstrates promising progress towards using TK-style architectures for retrieval from the full collection, it is worthwhile to highlight several challenges that needs further explorations. More in depth analysis of the distribution of term-document scores is necessary which may divulge further insights about how sparsity properties and discretization can be enforced for practical operationlization of these models. Large scale pretraining in the context of these models also presents itself as an important direction for future studies. Finally, for the full retrieval setting, identifying appropriate negative document sampling strategies during training poses as an important challenge that can strongly help or curtail the success these models achieve on these tasks.

In the first year of the TREC Deep Learning track, there was a stronger focus on the reranking setting—although some submissions explored document expansion and other QTI-based strategies. We anticipate that in the 2020 edition of the track, we will observe more submissions using neural methods for the full retrieval setting, which may further improve the reusability of the TREC benchmark [Yilmaz et al., 2020] for comparing these emerging family of approaches, and provide additional insights for our line of exploration.

References

- Amin Ahmad, Noah Constant, Yinfei Yang, and Daniel Cer. Reqa: An evaluation for end-to-end answer retrieval models. *arXiv preprint arXiv:1907.04780*, 2019.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proc. ICML*, pages 89–96. ACM, 2005.
- Wei-Cheng Chang, Felix X Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. Pre-training tasks for embedding-based large-scale retrieval. *arXiv preprint arXiv:2002.03932*, 2020.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *CoRR*, abs/1904.10509, 2019. URL <http://arxiv.org/abs/1904.10509>.
- Nick Craswell. Mean reciprocal rank. In *Encyclopedia of Database Systems*, pages 1703–1703. Springer, 2009.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. Overview of the trec 2019 deep learning track. In *Proc. TREC*, 2019.

- Nick Craswell, Daniel Campos, Bhaskar Mitra, Emine Yilmaz, and Bodo Billerbeck. Orcas: 18 million clicked query-document pairs for analyzing search. *arXiv preprint arXiv:2006.05324*, 2020.
- Zhuyun Dai and Jamie Callan. Context-aware passage term weighting for first stage retrieval.
- Zhuyun Dai and Jamie Callan. Deeper text understanding for ir with contextual neural language modeling. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 985–988, 2019a.
- Zhuyun Dai and Jamie Callan. An evaluation of weakly-supervised deepct in the trec 2019 deep learning track. In *TREC*, 2019b.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. Neural ranking models with weak supervision. In *Proc. SIGIR*, pages 65–74. ACM, 2017.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Jiafeng Guo, Yixing Fan, Liang Pang, Liu Yang, Qingyao Ai, Hamed Zamani, Chen Wu, W. Bruce Croft, and Xueqi Cheng. A deep look into neural ranking models for information retrieval. *Information Processing & Management*, 2019.
- Sebastian Hofstätter, Hamed Zamani, Bhaskar Mitra, Nick Craswell, and Allan Hanbury. Local self-attention over long text for efficient document retrieval. In *Proc. SIGIR*. ACM, 2020a.
- Sebastian Hofstätter, Markus Zlabinger, and Allan Hanbury. Interpretable & time-budget-constrained contextualization for re-ranking. In *Proc. of ECAI*, 2020b.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*, 2016.
- Vladimir Karpukhin, Barlas Öğuz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*, 2020.
- Omar Khattab and Matei Zaharia. Colbert: Efficient and effective passage search via contextualized late interaction over bert. *arXiv preprint arXiv:2004.12832*, 2020.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2019.
- Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, pages 489–504, 2018.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*, 2019.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. Sparse, dense, and attentional representations for text retrieval. *arXiv preprint arXiv:2005.00181*, 2020.
- Ji Ma, Ivan Korotkov, Yinfei Yang, Keith Hall, and Ryan McDonald. Zero-shot neural retrieval via domain-targeted synthetic query generation. *arXiv preprint arXiv:2004.14503*, 2020.
- Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonellotto, Nazli Goharian, and Ophir Frieder. Expansion via prediction of importance with contextualization. *arXiv preprint arXiv:2004.14245*, 2020.
- Joel Mackenzie, Zhuyun Dai, Luke Gallagher, and Jamie Callan. Efficiency implications of term weighting for passage retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2020.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Proc. NIPS*, pages 3111–3119, 2013b.

- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, pages 746–751. Citeseer, 2013c.
- Bhaskar Mitra and Nick Craswell. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 2018.
- Bhaskar Mitra and Nick Craswell. Duet at trec 2019 deep learning track. In *Proc. TREC*, 2019a.
- Bhaskar Mitra and Nick Craswell. An updated duet model for passage re-ranking. *arXiv preprint arXiv:1903.07666*, 2019b.
- Bhaskar Mitra, Eric Nalisnick, Nick Craswell, and Rich Caruana. A dual embedding space model for document ranking. *arXiv preprint arXiv:1602.01137*, 2016.
- Bhaskar Mitra, Fernando Diaz, and Nick Craswell. Learning to match using local and distributed representations of text for web search. In *Proc. WWW*, pages 1291–1299, 2017.
- Bhaskar Mitra, Corby Rosset, David Hawking, Nick Craswell, Fernando Diaz, and Emine Yilmaz. Incorporating query term independence assumption for efficient retrieval and ranking using deep neural networks (under review). In *Proc. ACL*, 2019.
- Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. Improving document ranking with dual word embeddings. In *Proc. WWW*, 2016.
- Federico Nanni, Bhaskar Mitra, Matt Magnusson, and Laura Dietz. Benchmark for complex answer retrieval. In *Proc. ICTIR*, pages 293–296. ACM, 2017.
- Rodrigo Nogueira and Kyunghyun Cho. Task-oriented query reformulation with reinforcement learning. In *Proc. EMNLP*, pages 574–583, 2017.
- Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*, 2019.
- Rodrigo Nogueira, Jimmy Lin, and AI Epistemic. From doc2query to docttttquery. *Online preprint*, 2019a.
- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. Document expansion by query prediction. *arXiv preprint arXiv:1904.08375*, 2019b.
- Harrie Oosterhuis, J Shane Culpepper, and Maarten de Rijke. The potential of learned index structures for index compression. In *Proceedings of the 23rd Australasian Document Computing Symposium*, pages 1–4, 2018.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Łukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. *arXiv preprint arXiv:1802.05751*, 2018.
- Stephen Robertson, Hugo Zaragoza, et al. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389, 2009.
- Corby Rosset, Damien Jose, Gargi Ghosh, Bhaskar Mitra, and Saurabh Tiwary. Optimizing query evaluations using reinforcement learning for web search. In *Proc. SIGIR*. ACM, 2018.
- Corby Rosset, Bhaskar Mitra, Chenyan Xiong, Nick Craswell, Xia Song, and Saurabh Tiwary. An axiomatic approach to regularizing neural ranking models. In *Proc. SIGIR*, pages 981–984, 2019.
- Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *arXiv preprint arXiv:2003.05997*, 2020.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using gpu model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. Adaptive attention span in transformers. *arXiv preprint arXiv:1905.07799*, 2019.
- Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. Sparse sinkhorn attention. *arXiv preprint arXiv:2002.11296*, 2020.
- Christophe Van Gysel, Bhaskar Mitra, Matteo Venanzi, Roy Rosemarin, Grzegorz Kukla, Piotr Grudzien, and Nicola Cancedda. Reply with: Proactive recommendation of email attachments. In *Proc. CIKM*, 2017.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.

- Lilian Weng. Attention? attention! *lilianweng.github.io/lil-log*, 2018. URL <https://lilianweng.github.io/lil-log/2018/06/24/attention-attention.html>.
- Lilian Weng. The transformer family. *lilianweng.github.io/lil-log*, 2020. URL <https://lilianweng.github.io/lil-log/2020/04/07/the-transformer-family.html>.
- Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. Lite transformer with long-short range attention. In *International Conference on Learning Representations*, ICLR ’20, 2020.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. *arXiv preprint arXiv:2007.00808*, 2020.
- Ming Yan, Chenliang Li, Chen Wu, Bin Bi, Wei Wang, Jiangnan Xia, and Luo Si. Idst at trec 2019 deep learning track: Deep cascade ranking with generation-based document expansion and pre-trained language modeling. In *TREC*, 2019.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
- Emine Yilmaz, Nick Craswell, Bhaskar Mitra, and Daniel Campos. On the reliability of test collections to evaluating systems of different types. In *Proc. SIGIR*. ACM, 2020.
- Zeynep Akkalyoncu Yilmaz, Shengjin Wang, and Jimmy Lin. H2oloo at trec 2019: Combining sentence and document evidence in the deep learning track. In *TREC*, 2019.
- Hamed Zamani, Mostafa Dehghani, W Bruce Croft, Erik Learned-Miller, and Jaap Kamps. From neural re-ranking to neural ranking: Learning a sparse representation for inverted indexing. In *Proc. CIKM*, pages 497–506. ACM, 2018.