

A Systematic Evaluation of Transfer Learning and Pseudo-labeling with BERT-based Ranking Models

Iurii Mokrii*
ymokriy@hse.ru
HSE University
Moscow, Russia

Leonid Boytsov*
leonid.boytsov@us.bosch.com
Bosch Center for Artificial
Intelligence
Pittsburgh, USA

Pavel Braslavski
Ural Federal University
Yekaterinburg, Russia
HSE University
Moscow, Russia

ABSTRACT

Due to high annotation costs making the best use of existing human-created training data is an important research direction. We, therefore, carry out a systematic evaluation of transferability of BERT-based neural ranking models across five English datasets. Previous studies focused primarily on zero-shot and few-shot transfer from a large dataset to a dataset with a small number of queries. In contrast, each of our collections has a substantial number of queries, which enables a full-shot evaluation mode and improves reliability of our results. Furthermore, since source datasets licences often prohibit commercial use, we compare transfer learning to training on pseudo-labels generated by a BM25 scorer. We find that training on pseudo-labels—possibly with subsequent fine-tuning using a modest number of annotated queries—can produce a competitive or better model compared to transfer learning. Yet, it is necessary to improve the stability and/or effectiveness of the few-shot training, which, sometimes, can degrade performance of a pretrained model.

CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking.**

KEYWORDS

Neural information retrieval, transfer learning, pseudo-labeling

ACM Reference Format:

Iurii Mokrii, Leonid Boytsov, and Pavel Braslavski. 2021. A Systematic Evaluation of Transfer Learning and Pseudo-labeling with BERT-based Ranking Models. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463093>

1 INTRODUCTION

A recent adoption of large pretrained Transformer models [9, 34] in information retrieval (IR) led to a substantially improvement of ranking accuracy compared to traditional, i.e., non-neural retrieval models [7]. It also enabled effective zero-shot transfer learning in a monolingual [1, 11, 28, 35, 36] and cross-lingual settings [1, 19, 29].

Transfer learning may reduce the need to collect expensive human relevance judgements required for supervised training [12, 15].

*Equal contribution.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in (SIGIR '21), *SIGIR '21, July 11–15, 2021, Virtual Event, Canada*.
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8037-9/21/07.
<https://doi.org/10.1145/3404835.3463093>

However, many source collections such as a popular large-scale MS MARCO [2] have a non-commercial, research-only license, which limits practical applicability of transfer learning. Furthermore, few-shot learning with transferred models may produce results inferior to transfer learning alone [36]. From the methodological point of view, prior studies focus primarily on zero-shot and few-shot transfer from a dataset with a large number of queries to a dataset with a small number of queries. We have also not seen a study that compares transfer learning to training of BERT-based models on pseudo-labels (generated using in-domain data) [8].

To fill the gap, we study transferability of BERT-based ranking models and compare transfer learning to training on pseudo-labels generated using a BM25 scoring function [27]. We use five diverse English datasets that differ in terms of document/query types and/or lengths. In contrast to previous studies, each of our collections has a substantial number of queries, which enables a full-shot evaluation mode and improves reliability of our results.

Importantly, this short paper focuses on evaluation of existing techniques rather than on improving them. We ask the following research questions:

- **RQ1:** When training from scratch, how much data does a BERT-based ranker need to outperform BM25?
- **RQ2:** Does a model trained on pseudo-labels outperform BM25 (and by how much)?
- **RQ3:** Is transfer learning always more effective than BM25?
- **RQ4:** Is transfer learning more effective than training on pseudo-labels?
- **RQ5:** Can we improve upon transfer learning and/or pseudo-labeling with a few training examples?

We find that:

- **RQ1:** Training a competitive BERT-based models from scratch may require a substantial number of annotated queries: From one hundred to several thousands.
- **RQ2:** Models trained only on pseudo-labels consistently outperform BM25 by 5-15%.
- **RQ3:** However, transferred models can be worse than BM25.
- **RQ4 and RQ5:** Transferred models are typically better than models trained on pseudo-labels, but we can often match or exceed performance of the former by training on a large number of pseudo-labels with subsequent fine-tuning using a moderate number of annotated queries.
- **RQ5:** In that, fine-tuning with a small number of annotated queries can cause a substantial performance degradation, which confirms prior findings [36].

Table 1: Effectiveness (MRR) of zero-shot (ZS) and full-shot (FS) transfer

Target → Source ↓	Yahoo! Answers		MS MARCO doc		MS MARCO pass		DPR NQ		DPR SQuAD	
	ZS	FS	ZS	FS	ZS	FS	ZS	FS	ZS	FS
Yahoo! Answers	–	0.32	0.28*	0.38	0.24	0.33	0.34	0.49	0.47*	0.65
MS MARCO doc	0.15*	0.33	–	0.38	0.29*	0.34	0.44*	0.51	0.56*	0.65
MS MARCO pass	0.19*	0.33	0.30	0.39	–	0.34	0.43*	0.49	0.52*	0.65
DPR NQ	0.25*	0.32	0.28*	0.38	0.24	0.33	–	0.49	0.53*	0.64
DPR SQuAD	0.24*	0.33	0.24*	0.38	0.22	0.33	0.40*	0.49	–	0.65
BM25	0.27		0.29		0.22		0.31		0.44	
pseudo-labelling	0.29	0.33	0.31[#]	0.38	0.23	0.32	0.35 [#]	0.48	0.50 [#]	0.64

Notes: Statistically significant differences between pseudo-labeling and BM25 are marked with #; statistically significant differences between transfer learning and pseudo-labeling are marked with *.

In summary, we find that pseudo-labeling (possibly combined with few-shot training) delivers results competitive with (and sometimes superior to) transfer learning. However, there is a need to improve the stability and/or effectiveness of the few-shot training.

2 RELATED WORK

A detailed discussion of neural ranking models can be found in recent surveys [10, 17, 21]. The success of early approaches was controversial [16], but the models relying on large pretrained Transformers [34], in particular, BERT-based models [9] decidedly outperformed prior neural and traditional models on a variety of *retrieval* tasks including TREC evaluations [7] and MS MARCO retrieval challenges.¹ Thus, BERT-based rankers are the focus of this study.

A number of studies demonstrated effectiveness of these models in zero-shot transfer learning. However, unlike our work, they do not explore fine-tuning in a large-data regime and use small test collections, which may affect reliability of results [3, 33].

Specifically, Yilmaz et al. [35] trained a BERT-based model on several collections with short passages and tested them on TREC newswire collections. By combining BM25 scores with the scores of a BERT-based ranker they outperformed prior approaches.

Rücklé et al. [28] analyzed transferability of BERT-based ranking models trained on questions from 140 StackExchange forums. They trained 140 models in a self-supervised fashion to retrieve question’s detailed description using a question title. They further evaluated 140 models on 9 external collections and found that BERT-based rankers outperformed traditional IR baselines in most cases.

Several studies employed zero-shot transfer in a cross-lingual setting. Shi and Lin [29] fine-tuned a multilingual BERT (mBERT) on the English TREC Microblog collection² and tested it on Chinese, Arabic, French, Hindi, and Bengali data (each having around 50 annotated topics). MacAvaney et al. [19] fine-tuned mBERT on TREC Robust04 and tested it on TREC newswire data in Arabic, Chinese, and Spanish (each featuring from 25 to 50 topics).

Although not directly related to our work, zero-shot transfer was evaluated with *traditional*, i.e., non-neural, learning-to-rank models. For example, Macdonald et al. found the transfer to be effective among different variants of ClueWeb collections [20].

¹<https://microsoft.github.io/msmarco/>

²<https://trec.nist.gov/data/microblog.html>

Table 2: Dataset statistics

Dataset	#queries train	#docs	#relev. /query	#tok. /query	#tok. /doc
Yahoo! Answers	100K	819.6K	5.7	11.9	63
MS MARCO doc	357K	3.2M	1	3.2	1197
MS MARCO pass	788.7K	8.8M	0.7	3.5	75
DPR NQ	53.9K	21M	7.9	4.5	141
DPR SQuAD	73.7K	21M	4.8	5	141

Notes: Development sets have 5K queries, test sets have 1.5K queries. Text length is the # of **BERT** word pieces.

However, not all studies demonstrated the superior transferability of BERT-based models compared to traditional IR baselines. Althammer et al. [1] experimented with BERT-based models trained on legal documents and zero-shot transferred them to a patent retrieval task. Transferred models were at par with BERT-based models trained on in-domain data, however, they were outperformed by a BM25 baseline. In the study of Thakur et al. [32] BM25 outperformed transferred BERT-based re-ranking models on six datasets out of 17. Similarly, in an answer-sentence retrieval task a BM25 scorer combined with BERT subword tokenization outperformed other methods on five out of eight datasets [11].

Several papers explored the relationship between the amount of training data and the effectiveness of the resulting IR model. In particular, Karpukhin et al. [14] showed that increasing the number of training examples gradually improved the quality of a passage retriever. Nogueira et al. [24] observed that T5 [25] significantly outperformed BERT in a data-poor regime. In these studies, using more training data always resulted in better performance. However, Zhang et al. [36] discovered that fine-tuning a BERT-based ranker with a few queries on TREC Robust04 collection led to a substantial degradation of performance compared to zero-shot transfer. This surprising result motivated our **RQ5**. One can train a neural ranker on pseudo-labels generated by a traditional retrieval model such as BM25 [27]. Although this approach had been shown to be successful in the past [8], we are not aware of any recent (and systematic) evaluation of pseudo-labeling with a BERT-based ranker.

3 DATA

We use five retrieval question-answering (QA) English datasets, whose statistics is summarized in Table 2. Our dataset selection rationale is twofold. First, we needed a large number of queries for evaluation in different regimes (from zero- to full-shot). This was particularly important to answer **RQ1**. Second, a representative evaluation requires collections that differ in terms of document/query type (e.g., Wikipedia, Web, community QA), query types (factoid vs. non-factoid), and query/document lengths.

The first dataset—Yahoo! Answers—is the community question answering (CQA) dataset, which has mostly non-factoid questions. Users of the service ask questions on virtually any topic while other community members provide answers. We use a high-quality subset of Yahoo! Answers created by Surdeanu et al. [31].³ We treat all available answers to a question as relevant documents, including answers that are not marked as “best answers”. Queries are created by concatenating short questions and their longer descriptions. We randomly split Yahoo! Answers into training, development, and testing subsets. We verify that the split has no obvious data leakage, i.e., that only a small fraction of the questions have duplicates or near-duplicates across splits.

MS MARCO document (MS MARCO doc) and passage (MS MARCO pass) retrieval collections are *related* datasets created from the MS MARCO reading comprehension dataset [2] and contain a large number of question-like queries sampled from the Bing search engine log with subsequent filtering. These queries are not necessarily proper English questions, e.g., “lyme disease symptoms mood”, but they are answerable by a short passage retrieved from a set of about 3.6M Web documents [2]. Relevance judgements are quite sparse (about one relevant passage/document per query) and a positive label indicates that the passage can answer the respective question. The document retrieval data set (MS MARCO doc) is created by transferring passage-level relevance to original documents from which passages were extracted [7]. Thus, a document is considered relevant only if it contains at least one relevant passage.

The DPR data sets were created by Karpukhin et al. [14] by matching Wikipedia passages with questions from two reading comprehension data sets: Natural Questions [15] and SQuAD v1.1 [26]; we denote respective datasets as DPR NQ and DPR SQuAD. They processed a Wikipedia dump by removing tables, infoboxes, etc., and split pages into 21M passages containing at most 100 words. We use relevance judgements, questions, and passages provided by the authors.⁴

All collections except Yahoo! Answers come with large “official” development sets containing at least 5K queries, a subset of which we used for testing. Hyper-parameter tuning was carried out on separate sets sampled from the original training data. For few- and medium-shot training, we randomly sampled training sets of progressively increasing sizes. Because we had to carry out a large number of experiments, we limited the number of samples to three per query set size and used only 1.5K randomly sampled test queries.

4 METHODS

We use a BM25 scorer [27] tuned on a development set as a main retrieval baseline. For each query, 100 documents with top BM25 scores are used as an input to a neural re-ranker as well as to create pseudo-labels [8]. Relevant pseudo-labels are created without human supervision by selecting a document with the highest BM25 score. We use all available training queries.

We use a 12-layer BERT_{BASE} [9, 34] with a fully-connected prediction layer as a neural re-ranker [23].⁵ BERT takes a query-document pair as an input. Long MS MARCO doc documents are truncated to 445 first BERT tokens, but such shortening leads to only small ($\approx 1\%$) loss in accuracy [4]. Likewise, we keep at most 64 BERT tokens in queries.

The models are trained using a pairwise margin loss (inference is pointwise). In a single training epoch, we select randomly one pair of positive and negative examples per query (negatives are sampled from 100 documents with highest BM25 scores). We use an AdamW [18] optimizer with a small weight decay (10^{-7}), a warm-up schedule and a batch size of 16.⁶ Note that we use different base rates for the fully-connected prediction head ($2 \cdot 10^{-4}$) and for the main Transformer layers ($2 \cdot 10^{-5}$).

We estimated a number of training epochs necessary to achieve good performance when training from scratch. To this end, we experimented with a small number of queries on a development set. We observe that for all collections, achieving good performance with only 32 queries required 16-32 epochs. We also observe that training for a larger number of epochs may lead to some overfitting, but the effect is quite small (1-3%). Thus, we start with 32 epochs for 32 queries and decrease the number of epochs as the training set size increases. We use this strategy for both training from scratch and fine-tuning a model.

Experiments are carried out using FlexNeuART [5] framework. Effectiveness is measured using the mean reciprocal rank (MRR), which is an official metric for MS MARCO data [7]. For statistical significance testing we use a paired t-test (threshold = 0.01).

5 CONCLUDING DISCUSSION OF RESULTS

Table 1 and Figure 1 contain experimental results. Figure 1 shows the relationship between the test accuracy and the training set size (measured in the number of queries). Because not all queries have relevant documents (especially in MS MARCO pass), these sizes are smaller than those in Table 2. Vertical bars indicate the range of test values for training samples of the same size. Different graphs in a panel correspond to training from a different starting model: There is graph for training from scratch, from a model trained on pseudo-labels, as well as one graph per each source collection.

RQ1: we can see that outperforming BM25 requires over 100 annotated queries on DPR data, at least 1-2K annotated queries on MS MARCO data and more than 8K annotated queries on Yahoo! Answers. Judging a single document-pair takes at least one minute on average [12, 15] and a single query typically needs at least 50 of such judgements [6]. Thus, annotating a single query by one

³Collection L6 in Yahoo WebScope: <https://webscope.sandbox.yahoo.com>

⁴<https://github.com/facebookresearch/DPR>

⁵BERT_{BASE} performs at par with BERT_{LARGE} on MS MARCO [13] and thus is a more practical alternative.

⁶The learning rate grows linearly from zero for 20% of the steps until it reaches the base learning rate [22, 30] and then goes back to zero (also linearly).

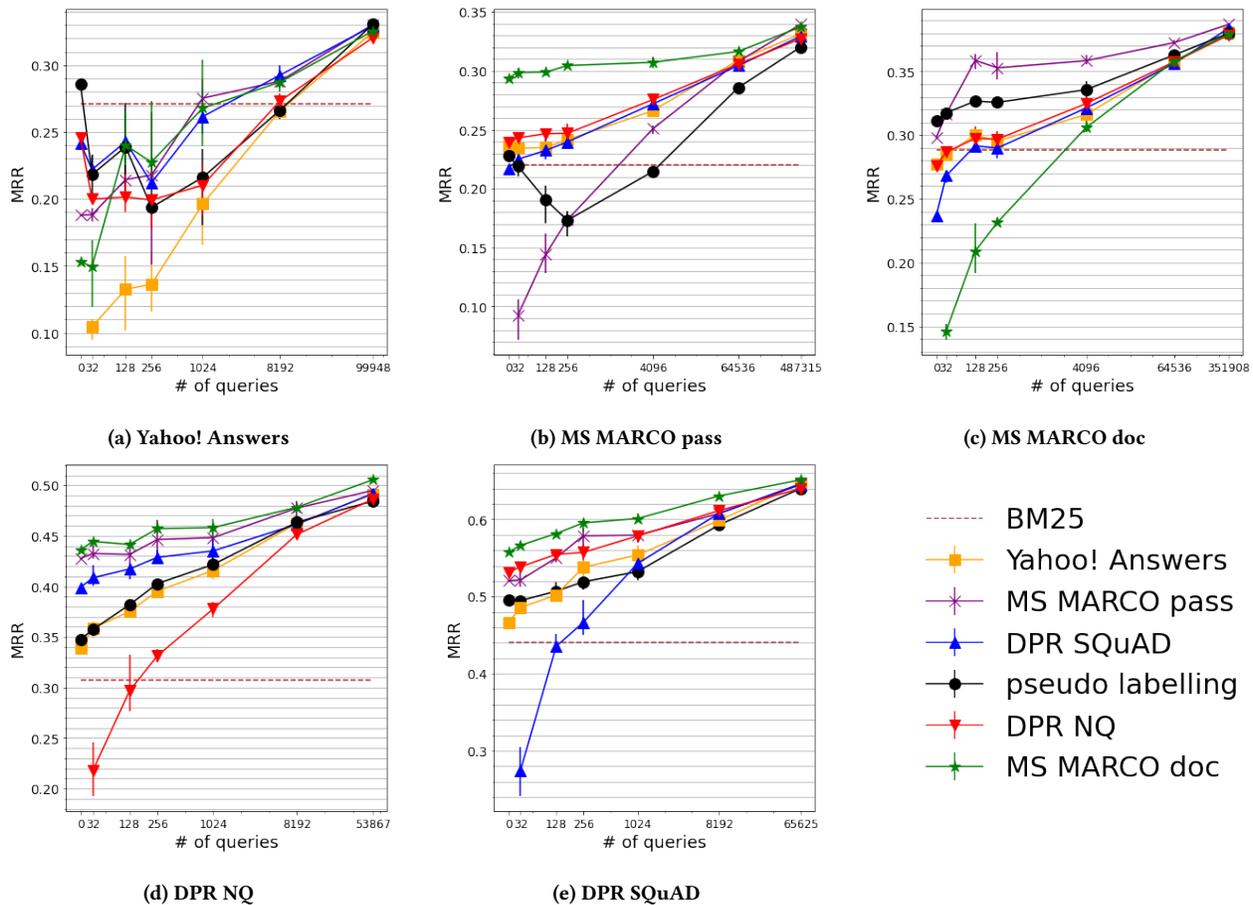


Figure 1: The relationship between MRR and the number of training queries.

worker can take longer than an hour. For MS MARCO, this entails several person-months of work just to match the accuracy of BM25.

RQ3: Transfer learning, however, can be worse than BM25 or outperform it only by a small margin (Table 1), which is in line with some prior work [1, 11, 32]. For example, for DPR collections a model transferred from Yahoo! Answers is only $\approx 10\%$ better than BM25. For Yahoo! Answers, *all* transferred models are worse than BM25. Transfer learning is also mostly *ineffective* on MS MARCO where only MS MARCO doc model transferred to a related MS MARCO pass dataset outperformed BM25 by as much as 30%.

RQ2: In contrast, pseudo-labeling *consistently* outperforms BM25 (differences are statistically significant except for Yahoo! Answers and MS MARCO pass). Yet, the observed gains (5-15%) are substantially smaller than those reported by Dehghani et al. [8].

RQ4 and RQ5: For almost every source model on DPR and MS MARCO datasets, a relatively small number of annotated queries (100-200) allow us to substantially improve upon both the transferred models and models trained on pseudo-labels. However, we also observe an “Little Bit Is Worse Than None” effect [36] on MS MARCO pass with pseudo-labeling as well on Yahoo! Answers.

The effect is particularly pronounced on Yahoo! Answers, where few-shot training ruins performance of *every* source model. We

hypothesize that few-shot training can lead to substantial overfitting to a small training set so that a model “forgets” what it learned from source training data.

We believe a similar forgetting happens when the amount of in-domain training data becomes sufficiently large (but it does not have a negative effect). As the number of training samples increases, the difference between different pretraining setups decreases: When we train using all the data, there is virtually no difference between starting from scratch or from a pretrained model.

To conclude, we note that transferred models are typically better than models trained on pseudo-labels and these differences are mostly statistically significant (see Table 1). However, we can often match or exceed performance of transferred models using a modest number of annotated queries to fine-tune a model trained on *pseudo-labels*. We thus, hypothesize, that training on pseudo-labels with a subsequent few-shot training on human-annotated data can become a viable alternative to transfer learning. Unlike zero-shot models trained on out-of-domain data, this scenario uses only in-domain data. Thus, it is likely to be less affected by the distribution mismatch between training and testing sets. However, one needs to improve the stability and effectiveness of the few-shot training, which, nevertheless, is out of the scope of this short paper.

ACKNOWLEDGMENTS

Pavel Braslavski thanks the Ministry of Science and Higher Education of the Russian Federation (“Ural Mathematical Center” project).

REFERENCES

- [1] Sophia Althammer, Sebastian Hofstätter, and Allan Hanbury. 2020. Cross-domain Retrieval in the Legal and Patent Domains: a Reproducibility Study. *arXiv preprint arXiv:2012.11405* (2020).
- [2] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. MS MARCO: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016).
- [3] Leonid Boytsov, Anna Belova, and Peter Westfall. 2013. Deciding on an adjustment for multiplicity in IR experiments. In *SIGIR*. 403–412.
- [4] Leonid Boytsov and Zico Kolter. 2021. Exploring Classic and Neural Lexical Translation Models for Information Retrieval: Interpretability, Effectiveness, and Efficiency Benefits. In *ECIR*. 63–78.
- [5] Leonid Boytsov and Eric Nyberg. 2020. Flexible retrieval with NMSLIB and FlexNeuART. In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*. 32–43.
- [6] Chris Buckley, Darrin Dimmick, Ian Soboroff, and Ellen M. Voorhees. 2007. Bias and the limits of pooling for large collections. *Inf. Retr.* 10, 6 (2007), 491–508.
- [7] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the TREC 2019 deep learning track. *arXiv preprint arXiv:2003.07820* (2020).
- [8] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural ranking models with weak supervision. In *SIGIR*. 65–74.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [10] Jiafeng Guo, Yixing Fan, Liang Pang, Liu Yang, Qingyao Ai, Hamed Zamani, Chen Wu, W Bruce Croft, and Xueqi Cheng. 2020. A deep look into neural ranking models for information retrieval. *Information Processing & Management* 57, 6 (2020), 102067.
- [11] Mandy Guo, Yinfei Yang, Daniel Cer, Qinlan Shen, and Noah Constant. 2020. MultiReQA: A Cross-Domain Evaluation for Retrieval Question Answering Models. *arXiv preprint arXiv:2005.02507* (2020).
- [12] Lei Han, Eddy Maddalena, Alessandro Checco, Cristina Sarasua, Ujwal Gadri-raj, Kevin Roitero, and Gianluca Demartini. 2020. Crowd Worker Strategies in Relevance Judgment Tasks. In *WSDM*. 241–249.
- [13] Sebastian Hofstätter, Markus Zlabinger, and Allan Hanbury. 2020. Interpretable & Time-Budget-Constrained Contextualization for Re-Ranking. In *ECAI*. 513–520.
- [14] Vladimir Karpukhin, Barlas Ögüz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. *arXiv preprint arXiv:2004.04906* (2020).
- [15] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural Questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics* 7 (2019), 453–466.
- [16] Jimmy Lin. 2019. The neural hype and comparisons against weak baselines. In *ACM SIGIR Forum*, Vol. 52. 40–51.
- [17] Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. 2020. Pretrained Transformers for Text Ranking: BERT and Beyond. *arXiv preprint arXiv:2010.06467* (2020).
- [18] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- [19] Sean MacAvaney, Luca Soldaini, and Nazli Goharian. 2020. Teaching a New Dog Old Tricks: Resurrecting Multilingual Retrieval Using Zero-Shot Learning. In *ECIR*. Springer, 246–254.
- [20] Craig Macdonald, Bekir Taner Dincer, and Iadh Ounis. 2015. Transferring Learning To Rank Models for Web Search. In *ICTIR*. ACM, 41–50.
- [21] Bhaskar Mitra and Nick Craswell. 2019. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* 13, 1 (2019), 1–126.
- [22] Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2020. On the Stability of Fine-tuning BERT: Misconceptions, Explanations, and Strong Baselines. *arXiv preprint arXiv:2006.04884* (2020).
- [23] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085* (2019).
- [24] Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. 2020. Document ranking with a pretrained sequence-to-sequence model. *arXiv preprint arXiv:2003.06713* (2020).
- [25] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683* (2019).
- [26] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100, 000+ Questions for Machine Comprehension of Text. In *EMNLP*. 2383–2392.
- [27] Stephen Robertson. 2004. Understanding inverse document frequency: on theoretical arguments for IDF. *Journal of Documentation* 60, 5 (2004), 503–520. <https://doi.org/10.1108/00220410410560582>
- [28] Andreas Rücklé, Jonas Pfeiffer, and Iryna Gurevych. 2020. MultiCQA: Zero-Shot Transfer of Self-Supervised Text Matching Models on a Massive Scale. *arXiv preprint arXiv:2010.00980* (2020).
- [29] Peng Shi and Jimmy Lin. 2019. Cross-lingual relevance transfer for document retrieval. *arXiv preprint arXiv:1911.02989* (2019).
- [30] Leslie N. Smith. 2017. Cyclical Learning Rates for Training Neural Networks. In *WACV*. 464–472.
- [31] Mihai Surdeanu, Massimiliano Ciaramita, and Hugo Zaragoza. 2011. Learning to rank answers to non-factoid questions from web collections. *Computational linguistics* 37, 2 (2011), 351–383.
- [32] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. *arXiv preprint arXiv:2104.08663* (2021).
- [33] Julián Urbano, Mónica Marrero, and Diego Martín. 2013. On the measurement of test collection reliability. In *SIGIR*. 393–402.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*. 5998–6008.
- [35] Zeynep Akkalyoncu Yilmaz, Wei Yang, Haotian Zhang, and Jimmy Lin. 2019. Cross-domain modeling of sentence-level evidence for document retrieval. In *EMNLP-IJCNLP*. 3481–3487.
- [36] Xinyu Zhang, Andrew Yates, and Jimmy Lin. 2020. A Little Bit Is Worse Than None: Ranking with Limited Training Data. In *Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing*. 107–112.