

Few-Shot Text Ranking with Meta Adapted Synthetic Weak Supervision

Si Sun¹, Yingzhuo Qian², Zhenghao Liu², Chenyan Xiong³,
Kaitao Zhang², Jie Bao¹, Zhiyuan Liu², Paul Bennett³

¹Department of Electronic Engineering, Tsinghua University, Beijing, China

²Department of Computer Science and Technology, Tsinghua University, Beijing, China

Institute for Artificial Intelligence, Tsinghua University, Beijing, China

Beijing National Research Center for Information Science and Technology, China

³Microsoft Research, Redmond, USA

{s-sun17, qyz17, liu-zh16, zkt18}@mails.tsinghua.edu.cn

{bao, liuzy}@tsinghua.edu.cn

{chenyan.xiong, Paul.N.Bennett}@microsoft.com

Abstract

The effectiveness of Neural Information Retrieval (Neu-IR) often depends on a large scale of in-domain relevance training signals, which are not always available in real-world ranking scenarios. To democratize the benefits of Neu-IR, this paper presents MetaAdaptRank, a domain adaptive learning method that generalizes Neu-IR models from label-rich source domains to few-shot target domains. Drawing on source-domain massive relevance supervision, MetaAdaptRank contrastively synthesizes a large number of weak supervision signals for target domains and meta-learns to reweight these synthetic “weak” data based on their benefits to the target-domain ranking accuracy of Neu-IR models. Experiments on three TREC benchmarks in the web, news, and biomedical domains show that MetaAdaptRank significantly improves the few-shot ranking accuracy of Neu-IR models. Further analyses indicate that MetaAdaptRank thrives from both its contrastive weak data synthesis and meta-reweighted data selection. The code and data of this paper can be obtained from <https://github.com/thunlp/MetaAdaptRank>.

1 Introduction

Text retrieval aims to rank documents to either directly satisfy users’ search needs or find textual information for later processing components, e.g., question answering (Chen et al., 2017) and fact verification (Liu et al., 2020). Neural information retrieval (Neu-IR) models have recently shown advanced results in many ranking scenarios where massive relevance labels or clickthrough data are available (Mitra et al., 2018; Craswell et al., 2020).

The flip side is that the “data-hungry” nature of Neu-IR models yields mixed results in few-shot ranking scenarios that suffer from the shortage of labeled data and implicit user feedback (Lin, 2019; Yang et al., 2019). On ranking benchmarks with

only hundreds of labeled queries, there have been debates about whether Neu-IR, even with billions of pre-trained parameters (Zhang et al., 2020a), really outperforms traditional IR techniques such as feature-based models and latent semantic indexing (Yang et al., 2019; Roberts et al., 2020). In fact, many real-world ranking scenarios are few-shot, e.g., tail web queries that innately lack large supervision (Downey et al., 2007), applications with strong privacy constraints like personal and enterprise search (Chirita et al., 2005; Hawking, 2004), and domains where labeling requires professional expertise such as biomedical and legal search (Roberts et al., 2020; Arora et al., 2018).

To broaden the benefits of Neu-IR to few-shot scenarios, we present an adaptive learning method MetaAdaptRank that meta-learns to adapt Neu-IR models to target domains with synthetic weak supervision. For synthesizing weak supervision, we take inspiration from the work (Ma et al., 2021) that generates related queries for unlabeled documents in a zero-shot way, but we generate discriminative queries based on contrastive pairs of relevant (positive) and irrelevant (negative) documents. By introducing the negative contrast, MetaAdaptRank can subtly capture the difference between documents to synthesize more ranking-aware weak supervision signals. Given that synthetic weak supervision inevitably contains noises, MetaAdaptRank meta-learns to reweight these synthetic weak data and trains Neu-IR models to achieve the best accuracy on a small volume of target data. In this way, neural rankers can distinguish more useful synthetic weak supervision based on the similarity of the gradient directions of synthetic data and target data (Ren et al., 2018) instead of manual heuristics or trial-and-error data selection (Zhang et al., 2020b).

We conduct experiments on three TREC benchmarks, ClueWeb09, Robust04, and TREC-COVID, which come from the web, news, and biomedical

cal domains, respectively. MetaAdaptRank significantly improves the few-shot ranking accuracy of Neu-IR models across all benchmarks. We also empirically indicate that both contrastive weak data synthesis and meta-reweighted data selection contribute to MetaAdaptRank’s effectiveness. Compared to prior work (Ma et al., 2021; Zhang et al., 2020b), MetaAdaptRank not only synthesizes more informative queries and effective weak relevance signals but customizes more diverse and fine-grained weights on synthetic weak data to better adapt neural rankers to target few-shot domains.

2 Related Work

Recent Neu-IR methods have achieved promising results in modeling relevance matching patterns between queries and documents (Guo et al., 2016; Hui et al., 2017; Mitra et al., 2018). They have been extensively employed in ad-hoc text retrieval (Xiong et al., 2017b; Dai et al., 2018; Nogueira and Cho, 2019; Xiong et al., 2021) and later natural language processing (NLP) tasks (Lee et al., 2019; Liu et al., 2020; Qu et al., 2020).

The effectiveness of Neu-IR methods heavily relies on the end-to-end training with a large number of relevance supervision signals, e.g., relevance labels or user clicks. Nevertheless, such supervision signals are often insufficient in many ranking scenarios. The less availability of relevance supervision pushes some Neu-IR methods to freeze their embeddings to avoid overfitting (Yates et al., 2020). The powerful deep pre-trained language models, such as BERT (Devlin et al., 2019), also do not effectively alleviate the dependence of Neu-IR on a large scale of relevance training signals. Recent research even observes that BERT-based neural rankers might require more training data than shallow neural ranking models (Hofstätter et al., 2020; Craswell et al., 2020). Moreover, they may often be overly confident and more unstable in the learning process (Qiao et al., 2019).

A promising direction to alleviate the dependence of Neu-IR models on large-scale relevance supervision is to leverage weak supervision signals that are noisy but available at mass quantity (Zheng et al., 2019b; Dehghani et al., 2017; Yu et al., 2020). Through IR history, various weak supervision sources have been used to approximate query-document relevance signals, e.g., pseudo relevance labels generated by unsupervised retrieval methods (Dehghani et al., 2017; Zheng et al., 2019b),

and title-document pairs (MacAvaney et al., 2019). Recently, Zhang et al. (2020b) treat paired anchor texts and linked pages as weak relevance signals and propose a reinforcement-based data selection method ReInfoSelect, which learns to filter noisy anchor signals with trial-and-error policy gradients. Despite their convincing results, anchor signals are only available in web domains. Directly applying them to non-web domains may suffer from suboptimal outcomes due to domain gaps. To obtain weak supervision that adapts arbitrary domains, Ma et al. (2021) present a synthetic query generation method, which can be trained with source-domain relevance signals and applied on target-domain documents to generate related queries.

More recently, a novel meta-learning technique has shown encouraging progress on solving data noises and label biases in computer vision (Ren et al., 2018; Shu et al., 2019; Zheng et al., 2019a) and some NLP tasks (Zheng et al., 2019a; Wang et al., 2020b). To the best of our knowledge, this novel technique has not been well utilized in information retrieval and synthetic supervision settings.

3 Methodology

This section first recaps the preliminary of Neu-IR and then introduces our proposed MetaAdaptRank. The framework of our method is shown in Figure 1.

3.1 Preliminary of Neu-IR

The ad-hoc retrieval task is to calculate a ranking score $f(q, d; \theta)$ for a query q and a document d from a document set. In Neu-IR, the ranking score $f(\cdot; \theta)$ is calculated by a neural model, e.g., BERT, with parameters θ . The query q and the document d are encoded to the token-level representations H :

$$H = \text{BERT}([\text{CLS}] \circ q \circ [\text{SEP}] \circ d \circ [\text{SEP}]), \quad (1)$$

where \circ represents the concatenation operation. [CLS] and [SEP] are special tokens. The first token (“[CLS]”) representation H_0 is regarded as the representation of the q - d pair. Then the ranking score $f(q, d; \theta)$ of the pair can be calculated as:

$$f(q, d; \theta) = \tanh(\text{Linear}(H_0)). \quad (2)$$

The standard learning to rank loss $l_i(\theta)$ (Liu, 2009), e.g., pairwise loss, can be used to optimize the neural model with relevance supervision signals $\{(q_i, d_i^+, d_i^-), 1 \leq i \leq M\}$:

$$l_i(\theta) = \text{relu}(1 - (f(q_i, d_i^+; \theta) - f(q_i, d_i^-; \theta))), \quad (3)$$

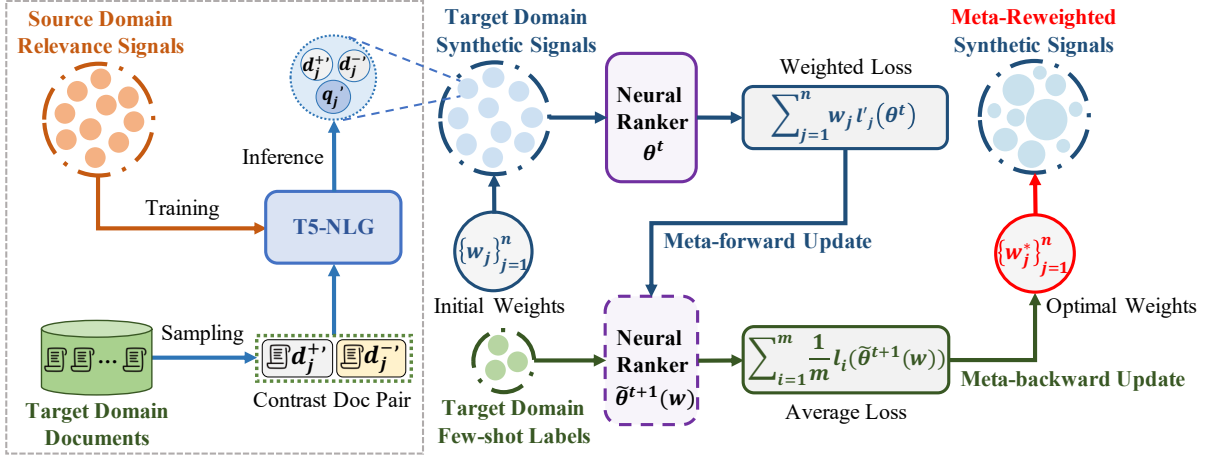


Figure 1: The illustration of MetaAdaptRank, which first synthesizes massive weak supervision signals for target domains, and then meta-learns to reweight these synthetic data based on small target-domain relevance labels.

where d_i^+ and d_i^- denote the relevant (positive) and irrelevant (negative) documents of the query q_i . In few-shot ranking scenarios, the number of relevance supervision signals (M) is limited, making it difficult to train an accurate Neu-IR model.

To mitigate the few-shot challenge in Neu-IR, MetaAdaptRank first transfers source-domain supervision signals to target-domain weak supervision signals (Sec 3.2); then meta-learns to reweight the synthetic weak supervision (Sec 3.3) for selectively training Neu-IR models (Sec 3.4).

3.2 Contrastive Synthetic Supervision

MetaAdaptRank transfers the relevance supervision signals from source domains to few-shot target domains in a zero-shot way. In this way, a natural language generation (NLG) model is trained on source domain relevance signals (*Source-domain NLG Training*) and is employed in target domains to synthesize weak supervision signals (*Target-domain NLG Inference*). We will first recap the previous synthetic method (Ma et al., 2021) and then introduce our contrastive synthetic approach.

Preliminary of Synthetic Supervision. Given a large volume of source-domain relevance pairs (q, d^+), previous synthetic method (Ma et al., 2021) trains a NLG model such as T5 (Raffel et al., 2020) that learns to generate a query q based on its relevant document d^+ :

$$q = \text{T5-NLG}([\text{POS}] \circ d^+ \circ [\text{SEP}]), \quad (4)$$

where [POS] and [SEP] are special tokens. In inference, the trained query generator is directly used to generate new queries q^* for target-domain documents d^* , where d^* is regarded as the related (posi-

tive) document of q^* , while the unrelated (negative) document can be sampled from the target corpus.

Despite some promising results, the vanilla training strategy may cause the NLG model to prefer to generate broad and general queries that are likely related to a crowd of documents in the target corpus. As a consequence, the synthetic relevance supervision does not have enough ranking awareness to train robust Neu-IR models.

Source-domain NLG Training. To synthesize ranking-aware weak supervision, MetaAdaptRank trains the NLG model to capture the difference between the contrastive document pair (d^+, d^-) and generate a discriminative query q :

$$q = \text{T5-NLG}([\text{POS}] \circ d^+ \circ [\text{NEG}] \circ d^- \circ [\text{SEP}]), \quad (5)$$

where [NEG] is another special token. The training instances (q, d^+, d^-) can be obtained from source domains in which d^+ and d^- are annotated as the relevant and irrelevant documents for the query q .

Target-domain NLG Inference. During inference, we first pick out a mass of confusable document pairs from target domains and then feed them into our trained contrastive query generator (Eq. 5) to synthesize more valuable weak supervision data.

To get confusable document pairs, we first generate a seed query q^* for each target-domain document d^* using the trained query generator (Eq. 4). Then the seed query is used to retrieve a subset of documents with BM25, where other retrieval methods can also be utilized. The confusable document pairs ($d^{+'}, d^{-'}$) are pairwise sampled from the retrieved subset without considering their rankings. Given the confusable document pair, we leverage our trained contrastive query generator to generate

a new query q' :

$$q' = \text{T5-NLG}([\text{POS}] \circ d^{+'} \circ [\text{NEG}] \circ d^{-'} \circ [\text{SEP}]), \quad (6)$$

where $d^{+'}$ and $d^{-'}$ are regarded as the related (positive) and unrelated (negative) documents of q' . In this way, we can synthesize massive target-domain weak supervision $\{(q_j', d_j^{+'}, d_j^{-'}), 1 \leq j \leq N\}$.

3.3 Meta Learning to Reweight

The synthetic weak data inevitably contain noises. To distinguish more useful training data for neural rankers, MetaAdaptRank meta-learns to reweight these synthetic data, following Ren et al. (2018).

Meta Learning Objective. Given a large volume of synthetic data $\{(q_j', d_j^{+'}, d_j^{-'}), 1 \leq j \leq N\}$ and a handful of target data $\{(q_i, d_i^+, d_i^-), 1 \leq i \leq M\}$ ($M \ll N$), our meta-learning objective is to find the optimal weights w^* on synthetic data to better train neural rankers. The learning of w^* involves *two nested loops of optimization*: initial-weighted synthetic data is used to *pseudo-optimize* the neural ranker; the weights is then *optimized* by minimizing the neural ranking loss on target data.

To be specific, the first loop (*Meta-forward Update*) incorporates the initial weights w into the learning parameters $\tilde{\theta}(w)$ instead of truly optimizing the neural ranker:

$$\tilde{\theta}(w) = \arg \min_{\theta} \sum_{j=1}^N w_j l_j'(\theta), \quad (7)$$

where $l_j'(\theta)$ is the ranking loss on a synthetic instance $(q_j', d_j^{+'}, d_j^{-'})$. In the second loop (*Meta-backward Update*), the optimal weights w^* can be obtained by minimizing the target ranking loss:

$$w^* = \arg \min_w \sum_{i=1}^M l_i(\tilde{\theta}(w)), \quad (8)$$

where $l_i(\theta)$ is the ranking loss on a target instance (q_i, d_i^+, d_i^-) . The calculation of each loop can be very expensive. In practice, we only perform one-step optimization in the two loops with mini-batch data, consistent with prior work (Ren et al., 2018).

Meta-forward Update. Taking the t -th training step as an example, we first assign a set of initial weights $w = \{w_j\}_{j=1}^n$ to the synthetic training data batch and then pseudo-update the neural ranker’s parameters to $\tilde{\theta}^{t+1}(w)$:

$$\tilde{\theta}^{t+1}(w) = \theta^t - \alpha \frac{\partial}{\partial(\theta^t)} \sum_{j=1}^n w_j l_j'(\theta^t), \quad (9)$$

where α is the learning rate. The description here uses vanilla SGD and other optimizers can be used.

Meta-backward Update. We leverage the neural ranker $\tilde{\theta}^{t+1}(w)$ to calculate the ranking loss on the target data batch and obtain the optimal weights $w^* = \{w_j^*\}_{j=1}^n$ through a single optimization step:

$$w_j^* = w_j - \eta \frac{\partial}{\partial(w_j)} \sum_{i=1}^m \frac{1}{m} l_i(\tilde{\theta}^{t+1}(w)), \quad (10)$$

where η is the learning rate for optimizing weights. The weights are further normalized for stable training. More details are shown in Appendices A.1.

3.4 Training with Meta-Weights

After obtaining the optimal weights w^* , the optimization of the neural ranker is a standard back-propagation on the weighted loss of synthetic data:

$$\theta^{t+1} = \arg \min_{\theta^t} \sum_{j=1}^n w_j^* l_j'(\theta^t). \quad (11)$$

In each training step, MetaAdaptRank first learns to reweight the synthetic batch based on their meta-impact on the target batch and then updates the neural ranker with the weighted synthetic batch. In this way, the few-shot target data can serve more as a “regularizer” to help the neural ranker to generalize with synthetic data, instead of as direct supervision which requires more labels (Ren et al., 2018).

4 Experimental Methodology

This section describes our experimental settings and implementation details.

Datasets. As shown in Table 1, three standard TREC datasets with different domains are used in our experiments: ClueWeb09-B (Callan et al., 2009), Robust04 (Kwok et al., 2004), and TREC-COVID (Roberts et al., 2020). They are all few-shot ad-hoc retrieval datasets where the number of labeled queries is limited. We leverage the “Complete” version of TREC-COVID whose retrieval document set is the July 16, 2020 release of COVID-19 (Wang et al., 2020a), a growing collection of scientific papers on COVID-19 and related research.

Evaluation Settings. We evaluate supervised IR methods through re-ranking the top 100 documents from the first-stage retrieval with five-fold cross-validation, consistent with prior work (Xiong et al., 2017a; Dai and Callan, 2019; Zhang et al., 2020b). The first-stage retrieval for ClueWeb09-B and Robust04 is the sequential dependence model

(SDM) (Metzler and Croft, 2005) released by Dai and Callan (2019), and the first-stage retrieval for TREC-COVID is BM25 (Robertson and Zaragoza, 2009) well-tuned by Anserini (Yang et al., 2017).

Metrics. NDCG@20 is used as the primary metric for all datasets. We also report ERR@20 for ClueWeb09-B and Robust04, which is the same with prior work (Zhang et al., 2020b), and report P@20 for TREC-COVID. Statistic significance is examined by permutation test with $p < 0.05$.

Baselines. Two groups of baselines are compared in our experiments, including *Traditional IR Baselines* and *Neural IR Baselines*.

Traditional IR Baselines. Following previous research (Dai and Callan, 2019; Zhang et al., 2020b), we compare four traditional IR methods in our experiments. They are two unsupervised methods, BM25 (Robertson and Zaragoza, 2009) and SDM (Metzler and Croft, 2005), and two learning-to-rank (LTR) methods using bag-of-word features, RankSVM (Joachims, 2002) and `Coor-Ascent` (Coordinate Ascent) (Metzler and Croft, 2007).

Neural IR Baselines. We also compare seven Neu-IR baselines that utilize different methodologies to train neural rankers. In our experiments, all Neu-IR methods adopt the widely-used BERT ranker (Nogueira and Cho, 2019), BERT-FirstP, which only uses the first paragraph of documents.

The vanilla neural baseline only leverages the existing small-scale relevance labels of target datasets to train BERT rankers, which is named `Few-shot Supervision`. We also compare BERT rankers trained with two large-scale supervision sources: `Bing User Click` and `MS MARCO`. Dai and Callan (2019) train BERT rankers with 5 million user click logs in Bing. We borrow their reported results because commercial logs are not publicly available. `MS MARCO` is a human supervision source (Nguyen et al., 2016), which provides over one million Bing queries with relevance labels.

Four weak supervision methods are also compared. One baseline is `Title Fitter`, which treats filtered title-document pairs as weak supervision signals (MacAvaney et al., 2019) for training BERT rankers (Zhang et al., 2020b). Another two baselines are `Anchor` and `ReInfoSelect`. `Anchor` leverages 100k pairs of anchor texts and web pages to train BERT rankers (Zhang et al., 2020b). `ReInfoSelect` first employs reinforcement learning to select these anchor signals (Zhang et al., 2020b) and then trains BERT rankers. The

Dataset	Domain	Corpus Size	Labeled Queries
ClueWeb09-B	Web Pages	50m	200
Robust04	News Articles	528k	250
TREC-COVID	BioMed Papers	191k	50

Table 1: Statistics of three TREC datasets used in our experiments. They are few-shot ranking datasets containing only tens to hundreds of labeled queries.

last baseline `SyncSup` trains BERT rankers with synthetic weak supervision data, which are synthesized based on the previous work (Ma et al., 2021).

Implementation Details. This part introduces the implement details of our method and baselines.

BERT Ranker. For our methods and all Neu-IR baselines, we use the base version of BERT (Devlin et al., 2019) on ClueWeb09-B and Robust04, and PubMedBERT (Base) (Gu et al., 2020) on TREC-COVID. We leverage the OpenMatch (Liu et al., 2021) implementation and obtain the pre-trained weights from Hugging Face (Wolf et al., 2020).

For all Neu-IR methods, we first use additional supervision sources such as weak supervision signals to train BERT rankers (except for `Few-shot Supervision`); then fine-tune the BERT rankers with the training folds of target datasets in the cross-validation. Following prior work (Dai and Callan, 2019; Zhang et al., 2020b), the ranking features ([CLS] embeddings) of BERT are combined with the first-stage retrieval scores using `Coor-Ascent` for ClueWeb09-B and Robust04. We set the max input length to 512 and use Adam optimizer with a learning rate of $2e-5$ and a batch size of 8.

Contrastive Supervision Synthesis. We use the small version of T5 (60 million parameters) as the NLG models in `MetaAdaptRank`, and leverage `MS MARCO` as the training data for T5-NLG models. We set the maximum input length to 512 and use Adam to optimize the T5-NLG models with a learning rate of $2e-5$ and a batch size of 4. In inference, the T5-NLG models are applied on target datasets with greedy search. Additionally, we consider `CTSyncSup` as our ablation baseline, which directly trains BERT rankers on contrastive synthetic supervision data without meta-reweighting.

Meta Learning to Reweight. The training folds of the target dataset are used as target data to guide the meta-reweighting to synthetic data. We set the batch size of synthetic data (n) and target data (m) to 8. The second-order gradient of the target ranking loss with regard to the initial weight (Eq. 10) is implemented using the automatic differentiation in PyTorch (Paszke et al., 2017).

Methods	ClueWeb09-B (Web)		Robust04 (News)		TREC-COVID (BioMed)	
	NDCG@20	ERR@20	NDCG@20	ERR@20	NDCG@20	P@20
BM25 (Yang et al., 2017)	0.2773	0.1426	0.4129	0.1117	0.6979	0.7670
SDM (Dai and Callan, 2019)	0.2774	0.1380	0.4269	0.1172	0.7030	0.7770
RankSVM (Dai and Callan, 2019)	0.289	n.a.	0.420	n.a.	n.a.	n.a.
RankSVM (OpenMatch)	0.2825	0.1476	0.4309	0.1173	0.6995	0.7570
Coor-Ascent (Dai and Callan, 2019)	0.295	n.a.	0.427	n.a.	n.a.	n.a.
Coor-Ascent (OpenMatch)	0.2969 [†]	0.1581 [†]	0.4340 [†]	0.1171	0.7041	0.7770
Few-shot Supervision (Zhang et al., 2020b)	0.2999	0.1631	0.4258	0.1163	n.a.	n.a.
Few-shot Supervision (Ours)	0.3033 [†]	0.1519	0.4572 ^{†‡}	0.1234	0.7713 ^{†‡}	0.8400 ^{†‡}
Bing User Click (Dai and Callan, 2019)	0.333	n.a.	n.a.	n.a.	n.a.	n.a.
MS MARCO (Nguyen et al., 2016)	0.3205 ^{†‡‡§}	0.1690 ^{†‡}	0.4674 ^{†‡}	0.1304 ^{†‡‡}	0.8054 ^{†‡‡}	0.8610 ^{†‡‡}
Title Filter (MacAvaney et al., 2019)	0.3021	0.1513	0.4379	0.1202	n.a.	n.a.
Anchor (Zhang et al., 2020b)	0.3072 [†]	0.1609 [†]	0.4449 ^{†‡}	0.1223 ^{†‡}	0.7677 ^{†‡}	0.8260 ^{†‡}
ReInfoSelect (Zhang et al., 2020b)	0.3261 ^{†‡‡§}	0.1669 ^{†‡}	0.4703 ^{†‡‡}	0.1313 ^{†‡‡}	0.7833 ^{†‡‡}	0.8420 ^{†‡‡}
SyncSup (Ma et al., 2021)	0.3036 [†]	0.1602 [†]	0.4685 ^{†‡}	0.1311 ^{†‡‡}	0.7867 ^{†‡}	0.8470 ^{†‡‡}
CTSyncSup	0.3123 [†]	0.1764 ^{†‡‡§}	0.4769 ^{†‡‡‡}	0.1293 ^{†‡‡}	0.8006 ^{†‡‡}	0.8610 ^{†‡‡}
MetaAdaptRank	0.3416 ^{†‡‡‡§}	0.1893 ^{†‡‡‡§}	0.4916 ^{†‡‡‡§}	0.1362 ^{†‡‡‡§}	0.8378 ^{†‡‡‡§}	0.8790 ^{†‡‡‡§}

Table 2: Ranking accuracy of MetaAdaptRank and baselines. †, ‡, †, ‡, ‡, ‡, ‡ indicate statistically significant improvements over SDM[†], Coor-Ascent[‡], Few-shot Supervision[†], MS MARCO[‡], ReInfoSelect[‡] and SyncSup[‡].

Supervision Sources	ClueWeb09-B (Web)		Robust04 (News)		TREC-COVID (BioMed)	
	NDCG@20	ERR@20	NDCG@20	ERR@20	NDCG@20	P@20
(a) MS MARCO (Nguyen et al., 2016)	0.3205 [†]	0.1690	0.4674 [‡]	0.1304 [‡]	0.8054 [‡]	0.8610 [‡]
(b) Anchor (Zhang et al., 2020b)	0.3072	0.1609	0.4449	0.1223	0.7677	0.8260
(c) SyncSup (Ma et al., 2021)	0.3036	0.1602	0.4685 [‡]	0.1311 [‡]	0.7867	0.8470
(d) CTSyncSup	0.3123	0.1764 [‡]	0.4769 [‡]	0.1293 [‡]	0.8006 [‡]	0.8610 [‡]
(e) MARCO + CTSyncSup	0.3214 [†]	0.1739 ^{†‡}	0.4727 [‡]	0.1297 [‡]	0.8182 ^{†‡}	0.8720 ^{†‡}

Table 3: Ranking accuracy with different supervision sources. MARCO + CTSyncSup denotes the hybrid source of MS MARCO and CTSyncSup. †, ‡, † indicate statistically significant improvements over (a)[†], (b)[‡] and (c)[†].

5 Evaluation Results

In this section, we present the evaluation results of MetaAdaptRank and conduct a series of analyses and case studies to study its effectiveness.

5.1 Overall Accuracy

The ranking results of MetaAdaptRank and baselines are presented in Table 2.

On all benchmarks and metrics, MetaAdaptRank outperforms all baselines stably. Compared to the best feature-based LeToR method, Coor-Ascent, MetaAdaptRank outperforms it by more than 15%. MetaAdaptRank even outperforms the strong Neu-IR baselines supervised with Bing User Click and MS MARCO, which demonstrates its effectiveness.

Specifically, CTSyncSup directly improves the few-shot ranking accuracy of BERT rankers by 3% on all benchmarks. In comparison to other weak supervision sources, filtered title-document relations, Anchor and SyncSup, CTSyncSup shows more stable effectiveness across different benchmarks, revealing its domain-adaptation advantages. Moreover, meta-reweighting CTSyncSup brings further improvement and helps MetaAdaptRank outperform the latest selective Neu-IR method ReInfoSelect.

Next, we go ahead to analyze MetaAdaptRank’s contrastive synthesis and meta-reweighting.

5.2 Effectiveness of Contrastive Synthesis

We analyze contrastive synthesis’s effectiveness by its effect on ranking results and synthetic quality.

Table 3 presents the ranking accuracy based on our CTSyncSup and four other supervision sources. CTSyncSup outperforms Anchor and SyncSup stably across all datasets. On Robust04, CTSyncSup even shows better performance than MS MARCO human labels. Besides, combining the sources of MS MARCO and CTSyncSup can further improve the ranking accuracy on ClueWeb09-B and TREC-COVID, revealing that CTSyncSup provides useful supervision signals applicable to various domains.

We further evaluate the quality of the queries generated in SyncSup and our CTSyncSup, which are both synthetic methods for generating queries based on target documents. Following previous research (Ma et al., 2021; Yu et al., 2020; Celikyilmaz et al., 2020), eight auto evaluation metrics are used in our evaluation. As shown in Table 4, CTSyncSup outperforms SyncSup on all metrics. The results demonstrate that the contrastive pair of positive and negative documents does help the

Synthetic Methods	BLEU-1	BLEU-2	ROUGE-1	ROUGE-2	ROUGE-L	NIST@1	NIST@2	METEOR
SyncSup (Ma et al., 2021)	0.5672	0.4527	0.5928	0.3764	0.5745	5.8070	7.3315	0.3089
Reverse-CTSyncSup	0.3185	0.1807	0.3528	0.1088	0.3395	3.0076	3.3665	0.1610
CTSyncSup	0.5909	0.4627	0.6238	0.3844	0.5955	6.1282	7.6314	0.3191

Table 4: Evaluation results of the queries generated by different synthetic methods. In Reverse-CTSyncSup, we swap the encoding order of contrastive document pairs, using original negative documents as positive documents.

Methods (Supervision Sources)	ClueWeb09-B (Web)		Robust04 (News)		TREC-COVID (BioMed)	
	NDCG@20	ERR@20	NDCG@20	ERR@20	NDCG@20	P@20
(a) ReInfoSelect (MS MARCO)	0.3294	0.1760	0.4756	0.1291	0.8229 [‡]	0.8780 [‡]
(b) ReInfoSelect (Anchor)	0.3261	0.1669	0.4703	0.1313	0.7891	0.8430
(c) ReInfoSelect (CTSyncSup)	0.3243	0.1742	0.4816 [‡]	0.1334	0.8230 [‡]	0.8800 [‡]
(d) MetaAdaptRank (MS MARCO)	0.3453 ^{†‡‡‡‡}	0.2018^{†‡‡‡‡}	0.4853 [‡]	0.1331	0.8354 ^{‡‡}	0.8730 [‡]
(e) MetaAdaptRank (Anchor)	0.3374	0.1730	0.4797	0.1314	0.8045	0.8650
(f) MetaAdaptRank (CTSyncSup)	0.3416 ^b	0.1893 ^{‡‡}	0.4916 ^{†‡‡}	0.1362 ^{†‡}	0.8378 ^{‡‡}	0.8790 [‡]
(g) MetaAdaptRank (MARCO + CTSyncSup)	0.3498^{†‡‡‡}	0.1926 ^{‡‡‡}	0.4989^{†‡‡‡‡}	0.1366^b	0.8488^{†‡‡‡‡}	0.8910^{†‡‡‡}

Table 5: Ranking accuracy of ReInfoSelect and MetaAdaptRank using different supervision sources. Superscripts †, ‡, †, ‡, †, ‡, †, ‡, †, ‡ indicate statistically significant improvements over (a)[†], (b)[‡], (c)[‡], (d)[‡], (e)[‡] and (f)[‡].

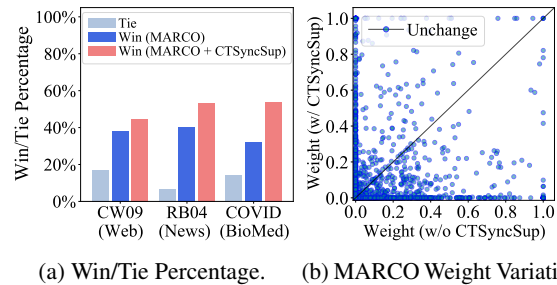
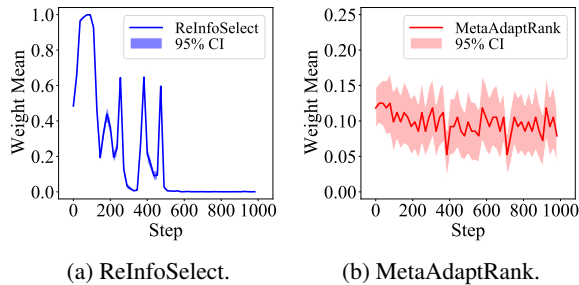


Figure 2: The state of learned weights on CTSyncSup data from ReInfoSelect and MetaAdaptRank. We use a ClueWeb09 few-shot fold as target data. Training steps are marked on X-Axes. The mean and 95% Confidence Interval (CI) of data weights in the same batch are plotted. A 95% CI is an interval that will contain the true mean of weights with 95% probability. Its width is proportional to the standard deviation of data weights.

Figure 3: The analysis of MetaAdaptRank on the hybrid source of MS MARCO and CTSyncSup. The ratio of Win/Tie queries between MS MARCO and the hybrid source is shown in (a). The statistics are based on NDCG@20 scores. CW09, RB04, COVID are short for datasets. (b) illustrates the variation in the meta-learned weights of 2k MS MARCO data points with (w/) and without (w/o) merging CTSyncSup.

NLG model better approximate the golden queries. In addition, reversing the encoding order of the contrastive document pair causes a dramatic decrease in all evaluation scores of the generated queries. This further shows that our contrastive query generator can extract more specific and representative information from the positive documents, thereby generating more discriminative queries.

5.3 Effectiveness of Meta Reweighting

To analyze the effectiveness of meta reweighting, we employ MetaAdaptRank on different supervision sources and study its data weighting behaviors in the learning process. The reinforcement data selector ReInfoSelect is used as a comparison, which utilizes the trial-and-error weighting mechanism.

The ranking accuracy of MetaAdaptRank and

ReInfoSelect trained with MS MARCO, Anchor, and CTSyncSup is presented in Table 5. For all supervision sources, MetaAdaptRank outperforms ReInfoSelect on all benchmarks. The results show that the meta-reweighting mechanism can more effectively explore the potential of different supervision sources compared to the trial-and-error weighting mechanism. Moreover, the advantages of meta reweighting can be extended to the hybrid supervision source of MS MARCO and CTSyncSup.

To further understand the behaviors of meta reweighting, we compare the state of weights assigned to synthetic supervision by MetaAdaptRank and ReInfoSelect in the learning process, using CTSyncSup as synthetic data and ClueWeb09 as target data. The results are shown in Figure 2. Even though each synthetic batch is likely to include both

TREC-COVID R5 Methods		All Queries		Old Queries		New Queries	
		NDCG@20	P@20	NDCG@20	P@20	NDCG@20	P@20
(a)	r5.fusion1 (Anserini BM25)	0.5313	0.5840	0.5202	0.5722	0.6320	0.6900
(b)	r5.fusion2 (Anserini BM25)	0.6007 [†]	0.6440 [†]	0.5937 [†]	0.6344 [†]	0.6641	0.7300
(c)	covidex.r5.2s (RRF)	0.7457 ^{†‡}	0.7610 ^{†‡}	0.7303 ^{†‡}	0.7456 ^{†‡}	0.8837 [†]	0.9000
(d)	MetaAdaptRank (rerank (a))	0.7536 ^{†‡}	0.7820 ^{†‡}	0.7405 ^{†‡}	0.7656 ^{†‡}	0.8712 ^{†‡}	0.9300 ^{†‡}
(e)	covidex.r5.d2q.2s (RRF)	0.7539 ^{†‡}	0.7700 ^{†‡}	0.7385 ^{†‡}	0.7544 ^{†‡}	0.8929 [†]	0.9100
(f)	MetaAdaptRank (rerank (b))	0.7904 ^{†‡‡‡}	0.8270 ^{†‡‡‡}	0.7790 ^{†‡‡‡}	0.8144 ^{†‡‡‡}	0.8933 ^{†‡‡}	0.9400 ^{†‡‡}
(g)	MetaAdaptRank (RRF)	0.7992 ^{†‡‡‡}	0.8380 ^{†‡‡‡}	0.7899 ^{†‡‡‡}	0.8267 ^{†‡‡‡}	0.8833 ^{†‡‡}	0.9400 ^{†‡‡}

Table 6: Evaluation results of TREC-COVID R5. All Queries denotes all queries in R5. Old and New Queries denote queries that have been judged or not in previous rounds (R1-R4). (a) and (b) are the first-stage retrieval of other methods in this table. (c) and (e) are R5’s top 2 automatic systems. (g) is the reciprocal rank fusion (RRF) of (d) and (f). †, ‡, †‡, †‡‡, †‡‡‡ indicate statistically significant improvements over (a)[†], (b)[‡], (c)^{†‡}, (d)^{†‡}, (e)^{†‡} and (f)^{†‡‡}.

	Synthetic Query	Positive Document	Negative Document
(↑)	CTSyncSup: how does shopping with the planet make a big difference in msn eco SyncSup: what is green energy ecosystem	... green at msn shopping shopping with the planet in mind can make a big difference by msn shopping msn green updated: energy saving solutions conserving energy reduces co2 emissions eco adventure tours energy star pledge donate resources boater guide marinas harbormasters green thumb ride share candle light dinner bright idea ...
(↓)	CTSyncSup: what is the history of bermuda SyncSup: where is <i>jamestown beach</i>	... bermuda beach resorts: website dedicated to advertising in bermuda large helpful travel forum bermuda links: activities hotels resorts <i>beach bermuda history bermuda</i> hotels your art history reference guide art history search <i>jamestown</i> , colonial history virginia (redirected from <i>jamestown</i> settlement) <i>jamestown</i> was a village on an island ...

Table 7: Cases of meta-reweighted contrastive synthetic data targeting ClueWeb09. The weights are marked in the parenthesis ↑ (more important) and ↓ (down-weight). The **red** texts are specific contents of positive documents and the **blue** texts are shared by both positive and negative documents. The document snippets are manually selected.

useful and noisy data points, ReInfoSelect always assigns very high weights at the beginning and discards almost all synthetic data points later. Besides, its tight confidence interval reveals that data points in the same batch received almost identical weights. These observations indicate that ReInfoSelect does not effectively distinguish useful synthetic data points from the noisy ones during the learning process. By contrast, MetaAdaptRank assigns higher weights initially and steadily reduces the weights as training goes on. More importantly, its wide confidence interval reveals that the data weights in the same synthetic batch vary significantly, which are thus expected to be more diverse and fine-grained.

5.4 Effectiveness of Hybrid Supervision

We also analyze MetaAdaptRank’s advantages on the hybrid supervision source of MS MARCO and CTSyncSup. The impact of the hybrid source on its ranking accuracy and meta-reweighting behavior is studied. Besides, we evaluate MetaAdaptRank trained with the hybrid source in Round 5 of the TREC-COVID shared task in which many strong baselines have been well-tuned for four rounds.

Figure 3a shows the Win/Tie ranking accuracy of MetaAdaptRank trained with MS MARCO and the

hybrid supervision source. Compared to the single MS MARCO, the hybrid source has more advantages across all benchmarks. Besides, the hybrid advantage seems to be more evident in non-web domain benchmarks, especially on TREC-COVID.

We further investigate the weighting behavior of MetaAdaptRank on MS MARCO and the hybrid source, using the same ClueWeb09 target data in previous analyses. Figure 3b illustrates the changes in meta-learned weights of randomly sampled 2k MS MARCO data points before and after merging CTSyncSup source. There are significant weight variations on most MS MARCO data points before and after merging CTSyncSup. Additionally, merging CTSyncSup reduces the weight of more MS MARCO data points, revealing that CTSyncSup data are assigned higher weights. This also reveals that MetaAdaptRank can tailor diversified weights for the same data points in different sources and up-weights more useful training data flexibly.

Lastly, we report the TREC-COVID R5 ranking results of MetaAdaptRank trained with the hybrid source. The top 2 automatic search systems in the R5 leaderboard are compared, which outperforms other systems on the newly added queries in R5. The evaluation of these new queries is fair to our

methods and those systems that underwent previous rounds (R1-R4). As shown in Table 6, our single model outperforms the top 2 fusion-based systems on all evaluation of the new, old, and all queries, further showing the effectiveness of MetaAdaptRank with the hybrid supervision source. More details and ranking results are shown in Appendices A.2.

5.5 Case Studies

Table 7 exhibits some cases of contrastive synthetic data for ClueWeb09 and their meta-learned weights. More cases are shown in Appendices A.3.

CTSyncSup can extract more specific contents from the positive documents, e.g., “shopping with the planet” and “make a big difference” in the first case; SyncSup captures more general information, e.g., “green energy”. Compared to SyncSup’s queries such as “where is jamestown beach” in the second case, the synthetic queries in CTSyncSup are more informative and discriminative. Noticeably, the second case exhibits the synthetic noise, where the positive document is actually related to “bermuda’s tourism” instead of the query “history of bermuda”. MetaAdaptRank effectively filters this noisy instance by assigning a zero weight to it.

6 Conclusion

This paper presents MetaAdaptRank, a domain adaption method for few-shot Neu-IR with contrastive weak data synthesis and meta-reweighted data selection. Contrastive synthesis generates informative queries and useful synthetic supervision signals. Meta-learned weights form high-resolution channels between target labels and synthetic signals, providing robust and fine-grained data selection for synthetic weak supervision. Both of them collaborate to significantly improve the neural ranking accuracy in various few-shot search scenarios.

Acknowledgments

This work is partly supported by the National Key Research and Development Program of China (No. 2020AAA0106501) and Beijing Academy of Artificial Intelligence (BAAI). We thank Zhuyun Dai and Jamie Callan for sharing the SDM results on ClueWeb09-B and Robust04 and thank Shi Yu for discussions in the query generation methodologies.

References

Piyush Arora, Murhaf Hossari, Alfredo Maldonado, Clare Conran, and Gareth JF Jones. 2018. *Chal-*

lenges in the development of effective systems for professional legal search. In *Proceedings of ProfS/KG4IR/Data: Search@ SIGIR*, pages 29–34.

Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. 2009. Clueweb09 data set.

Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. *Evaluation of text generation: A survey*. *arXiv preprint arXiv:2006.14799*.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. *Reading wikipedia to answer open-domain questions*. In *Proceedings of ACL*, pages 1870–1879.

Paul Alexandru Chirita, Wolfgang Nejdl, Raluca Paiu, and Christian Kohlschütter. 2005. *Using ODP meta-data to personalize search*. In *Proceedings of SIGIR*, pages 178–185.

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*.

Zhuyun Dai and Jamie Callan. 2019. *Deeper text understanding for IR with contextual neural language modeling*. In *Proceedings of SIGIR*, pages 985–988.

Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. *Convolutional neural networks for soft-matching n-grams in ad-hoc search*. In *Proceedings of WSDM*, pages 126–134.

Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. *Neural ranking models with weak supervision*. In *Proceedings of SIGIR*, pages 65–74.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Doug Downey, Susan Dumais, and Eric Horvitz. 2007. *Heads and tails: studies of web search with common and rare queries*. In *Proceedings of SIGIR*, pages 847–848.

Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. *Domain-specific language model pretraining for biomedical natural language processing*. *arXiv preprint arXiv:2007.15779*.

Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. *A deep relevance matching model for ad-hoc retrieval*. In *Proceedings of CIKM*, pages 55–64.

David Hawking. 2004. *Challenges in enterprise search*. In *Proceedings of ADC*, pages 15–24.

- Sebastian Hofstätter, Hamed Zamani, Bhaskar Mitra, Nick Craswell, and Allan Hanbury. 2020. [Local self-attention over long text for efficient document retrieval](#). In *Proceedings of SIGIR*, page 2021–2024.
- Zhiting Hu, Bowen Tan, Russ R Salakhutdinov, Tom M Mitchell, and Eric P Xing. 2019. [Learning data manipulation for augmentation and weighting](#). In *Proceedings of NeurIPS*, pages 15764–15775.
- Kai Hui, Andrew Yates, Klaus Berberich, and Gerard de Melo. 2017. [PACRR: A position-aware neural IR model for relevance matching](#). In *Proceedings of EMNLP*, pages 1049–1058.
- Thorsten Joachims. 2002. [Optimizing search engines using clickthrough data](#). In *Proceedings of KDD*, pages 133–142.
- Kui-Lam Kwok, Laszlo Grunfeld, HL Sun, Peter Deng, and N Dinstl. 2004. [TREC 2004 robust track experiments using PIRCS](#). In *Proceedings of TREC*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of ACL*, pages 6086–6096.
- Jimmy Lin. 2019. [The neural hype and comparisons against weak baselines](#). In *ACM SIGIR Forum*, volume 52, pages 40–51.
- Tie-Yan Liu. 2009. [Learning to rank for information retrieval](#). *Foundations and Trends in Information Retrieval*, 3(3):225–331.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. [Fine-grained fact verification with kernel graph attention network](#). In *Proceedings of ACL*, pages 7342–7351.
- Zhenghao Liu, Kaitao Zhang, Chenyan Xiong, Zhiyuan Liu, and Maosong Sun. 2021. [OpenMatch: An open source library for Neu-IR research](#). *arXiv preprint arXiv:2102.00166*.
- Ji Ma, Ivan Korotkov, Yinfei Yang, Keith Hall, and Ryan McDonald. 2021. [Zero-shot neural passage retrieval via domain-targeted synthetic question generation](#). In *Proceedings of EACL*, pages 1075–1088.
- Sean MacAvaney, Andrew Yates, Kai Hui, and Ophir Frieder. 2019. [Content-based weak supervision for ad-hoc re-ranking](#). In *Proceedings of SIGIR*, pages 993–996.
- Donald Metzler and W Bruce Croft. 2005. [A markov random field model for term dependencies](#). In *Proceedings of SIGIR*, pages 472–479.
- Donald Metzler and W Bruce Croft. 2007. [Linear feature-based models for information retrieval](#). *Information Retrieval*, 10(3):257–274.
- Bhaskar Mitra, Nick Craswell, et al. 2018. [An introduction to neural information retrieval](#). *Foundations and Trends® in Information Retrieval*, 13(1):1–126.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A human generated machine reading comprehension dataset](#). In *Proceedings of CoCo@ NIPS*, volume 1773.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. [Passage re-ranking with BERT](#). *arXiv preprint arXiv:1901.04085*.
- Joao Palotti, Harris Scells, and Guido Zuccon. 2019. [TrecTools: an open-source Python library for information retrieval practitioners involved in TREC-like campaigns](#). In *Proceedings of SIGIR*, pages 1325–1328.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. [Automatic differentiation in PyTorch](#). In *Proceedings of NeurIPS Autodiff Workshop*.
- Yifan Qiao, Chenyan Xiong, Zheng-Hao Liu, and Zhiyuan Liu. 2019. [Understanding the behaviors of BERT in ranking](#). *arXiv preprint arXiv:1904.07531*.
- Chen Qu, Liu Yang, Cen Chen, Minghui Qiu, W Bruce Croft, and Mohit Iyyer. 2020. [Open-retrieval conversational question answering](#). In *Proceedings of SIGIR*, pages 539–548.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Mengye Ren, Wenyan Zeng, Bin Yang, and Raquel Urtasun. 2018. [Learning to reweight examples for robust deep learning](#). In *Proceedings of ICML*, pages 4331–4340.
- Kirk Roberts, Tasmee Alam, Steven Bedrick, Dina Demner-Fushman, Kyle Lo, Ian Soboroff, Ellen Voorhees, Lucy Lu Wang, and William R Hersh. 2020. [TREC-COVID: rationale and structure of an information retrieval shared task for COVID-19](#). *Journal of the American Medical Informatics Association*, 27(9):1431–1436.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- G. Salton and Chris Buckley. 1997. [Improving retrieval performance by relevance feedback](#). *Journal of the Association for Information Science and Technology*, 41(4):355–364.
- Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. [Meta-WeightNet: learning an explicit mapping for sample weighting](#). In *Proceedings of NeurIPS*, pages 1919–1930.

- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, Kathryn Funk, Rodney Kinney, Ziyang Liu, William Merrill, et al. 2020a. [CORD-19: the covid-19 open research dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020b. [Balancing training for multilingual neural machine translation](#). In *Proceedings of ACL*, pages 8526–8537.
- Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of EMNLP*, pages 38–45.
- Chenyan Xiong, Jamie Callan, and Tie-Yan Liu. 2017a. [Word-entity duet representations for document ranking](#). In *Proceedings of SIGIR*, pages 763–772.
- Chenyan Xiong, Zhuyun Dai, Jamie Callan, Zhiyuan Liu, and Russell Power. 2017b. [End-to-end neural ad-hoc ranking with kernel pooling](#). In *Proceedings of SIGIR*, pages 55–64.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. [Approximate nearest neighbor negative contrastive learning for dense text retrieval](#). In *Proceedings of ICLR*.
- Peilin Yang, Hui Fang, and Jimmy Lin. 2017. [Anserini: Enabling the use of lucene for information retrieval research](#). In *Proceedings of SIGIR*, pages 1253–1256.
- Wei Yang, Kuang Lu, Peilin Yang, and Jimmy Lin. 2019. [Critically examining the "neural hype": Weak baselines and the additivity of effectiveness gains from neural ranking models](#). In *Proceedings of SIGIR*, pages 1129–1132.
- Andrew Yates, Siddhant Arora, Xinyu Zhang, Wei Yang, Kevin Martin Jose, and Jimmy Lin. 2020. [Capreolus: A toolkit for end-to-end neural ad hoc retrieval](#). In *Proceedings of WSDM*, pages 861–864.
- Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. [Few-shot generative conversational query rewriting](#). *Proceedings of SIGIR*, page 1933–1936.
- Edwin Zhang, Nikhil Gupta, Rodrigo Nogueira, Kyunghyun Cho, and Jimmy Lin. 2020a. [Rapidly deploying a neural search engine for the COVID-19 open research dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*.
- Kaitao Zhang, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. 2020b. [Selective weak supervision for neural information retrieval](#). In *Proceedings of WWW*, pages 474–485.
- Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. 2019a. [Meta label correction for learning with weak supervision](#). *arXiv preprint arXiv:1911.03809*.
- Yukun Zheng, Yiqun Liu, Zhi-Qiang Fan, Cheng Luo, Qingyao Ai, Min Zhang, and Shaoping Ma. 2019b. [Investigating weak supervision in deep ranking](#). *Data and Information Management*, 3:155–164.

Methods/Run ID	NDCG@20	P@20
UPrrf102-r5	0.7873	0.8000
UPrrf93-r5	0.7967	0.8200
covidex.r5.1s.lr	0.8019	0.8300
elhuyar_prf_nof99d	0.8209	0.8700
covidex.r5.d2q.1s.lr	0.8287	0.8400
elhuyar_prf_nof99p	0.8333	0.9000
MetaAdaptRank (rerank fusion.1)	0.8712	0.9300
UPrrf102-wt-r5	0.8804	0.9100
MetaAdaptRank (RRF)	0.8833	0.9400
covidex.r5.2s.lr	0.8837	0.9000
UPrrf93-wt-r5	0.8849	0.9100
covidex.r5.d2q.2s.lr	0.8929	0.9100
MetaAdaptRank (rerank fusion.2)	0.8933	0.9400

Table 8: Ranking results of our methods and baselines on the new queries of TREC-COVID R5. The baselines are the top 10 feedback systems in the R5 leaderboard, marked with their submitted ID. The three variants of MetaAdaptRank are the same as those in Table 6.

A Appendices

A.1 Batch Normalization of Meta-Weights

This part elaborates the batch normalization process for meta-learned weights. Following prior research (Ren et al., 2018), we first set the initial weights w to zeros and obtain the new weights \tilde{w} :

$$\tilde{w}_j = -\eta \frac{\partial}{\partial(w_j)} \sum_{i=1}^m \frac{1}{m} l_i(\tilde{\theta}^{t+1}(w)) \Big|_{w_j=0}. \quad (12)$$

Then we clip \tilde{w} to get non-negative weights \hat{w} and further normalize them in the batch to obtain the final weights w^* :

$$\begin{aligned} \hat{w}_j &= \max(0, \tilde{w}_j), \\ w_j^* &= \frac{\hat{w}_j}{(\sum_{p=1}^n \hat{w}_p) + \delta(\sum_{p=1}^n \hat{w}_p)}. \end{aligned} \quad (13)$$

Here $\delta(\sum_{p=1}^n \hat{w}_p) = 1$ when $\sum_{p=1}^n \hat{w}_p = 0$, to prevent division errors, otherwise it is 0. With the batch-normalization process, the hyperparameter η can be effectively eliminated. The normalization method is not constrained and other approaches can also be used (Shu et al., 2019; Hu et al., 2019).

A.2 Supplementary Results of TREC-COVID R5

This part supplements our evaluation results in the TREC-COVID R5 shared task. We will first recap the shared task and then present more evaluation results and our implementation details.

TREC-COVID R5. The TREC-COVID Challenge is an ad-hoc ranking task for COVID-19 literature, consisting of five rounds. TREC-COVID R5 is the last round of this challenge, where the document set is the July 16, 2020 version of CORD-19,

and the query set contains 50 testing queries. The first 45 queries have been used in previous rounds (R1-R4), and the last five queries are newly added in R5. As in previous rounds, TREC-COVID R5 adopts residual collection evaluation (Salton and Buckley, 1997). In residual collection evaluation, the relevance labels from previous rounds can be used, but any document that has been annotated for a query will be removed before the evaluation. We focus more on the evaluation of R5’s new queries because these queries have no prior relevance labels, which is fairer to our models and those search systems that underwent previous rounds.

Evaluation Results. Table 8 shows the evaluation results on the new queries of TREC-COVID R5, including three variants of our MetaAdaptRank and the top 10 feedback systems in the R5 leaderboard. Compared with the top 10 feedback systems (many are fusion-based systems), our single model MetaAdaptRank (rerank fusion.2) outperforms all baselines, demonstrating the generalization ability of our method on new queries.

Additionally, what catches our attention is that the best and worst of the top 10 feedback systems only have a 5.1% difference in NDCG@20 scores on all queries, while their NDCG@20 scores on the new queries differ by 13.4%. This discrepancy indicates that the residual collection evaluation may have biases between the seen and unseen queries.

Implementation Details. We next describe the implementation details of the three variants of our MetaAdaptRank in TREC-COVID R5. Consistent with the implementation methods described in Section 4, we rerank the top 100 documents from the first-stage retrieval. We first borrow two retrieval results with different settings provided by Anserini BM25 (Row 7 and 8 of Table Round 5¹). Then PubMedBERT (Base) is used to rerank these two retrieval results to obtain MetaAdaptRank (rerank fusion.1) and MetaAdaptRank (rerank fusion.2), respectively. MetaAdaptRank (RRF) is the reciprocal rank fusion of these two models. We utilize the open-source library trec-tools (Palotti et al., 2019) to implement RRF and set the fusion weight k to 1.

To train MetaAdaptRank, we first synthesize CTSyncSup data based on R5’s document set and leverage the hybrid source of CTSyncSup and MS MARCO as the additional supervision signals. The training process contains two stages. We first train

¹<https://github.com/castorini/anserini/blob/master/docs/experiments-covid.md>

	Synthetic Query	Positive Document	Negative Document
1 (↑)	CTSyncSup: us military radars in colombia SyncSup: what is the pentagon	... one month ago, the pentagon issued an order to suspend operations of the two radars that detect aircraft. these radars operate in colombia as a result of that agreement provide for more funding and retain more forces than the \$1.5-trillion five-year budget cheny presented to congress in january, pentagon officials say ...
2 (↑)	CTSyncSup: what percent of the economy was increased in 1993 SyncSup: what is the economic issue in peru	... this letter explains the peruvian government’s economic policy. the development of the economy in 1993 was in general much better. it is estimated that the real gdp has increased by 7 percent only three economies - guyana, argentina and peru - grew by more than than 5 per cent this year, with peru expanding by 11 percent ...
3 (↓)	CTSyncSup: what language is osvaldo rodriguez SyncSup: what is economic impact of cuba	... program with host osvaldo rodriguez the dialogue that was proposed by opponents to the revolutionary project. well, the word dialogue, is one that connotes cordiality. it is a positive word. but in the current political language , the counterrevolution’s political language program with host juan carlos roque garcia , and osvaldo rodriguez . this program could not ignore cuba’s presentation of a document by the u.s. interests section in havana to the un human rights commission in geneva ...
4 (↑)	CTSyncSup: which receptors are expressed in the human lung SyncSup: what is sars cov receptor	... results both sars-cov receptors of ace2 and cd209l were expressed in the 8 organ/tissue-derived endothelial cells. the expression of ace2 receptor was the highest in the human lung microvascular endothelial cells, and lowest 2019 novel coronavirus (2019- ncov) the outbreaks of 2002/2003 sars , 2012/2015 mers and 2019/2020 wuhan respiratory syndrome clearly indicate that genome evolution of an animal coronavirus (cov) may enable ...
5 (↑)	CTSyncSup: how does quarantine prevent covid outbreak SyncSup: covid outbreak symptoms	... the importance of the timing of quarantine measures before symptom onset to prevent covid-19 outbreaks how quarantine -based measures can prevent or suppress an outbreak furthermore, the effect of infectiousness prior to symptom onset combined with a significant proportion we evaluate two procedures: monitoring individuals for symptoms onset ...
6 (↓)	CTSyncSup: covid-19 pandemic effects on society SyncSup: what is the antiasia <i>sentiment</i> in the united states	... examination of community <i>sentiment</i> dynamics due to covid-19 pandemic : the outbreak of covid-19 has caused unprecedented impacts to people’s daily life around the world. virus may cause different mental health issues to people such as depression, anxiety, sadness mood of india during covid-19 - an interactive web portal based on emotion analysis of twitter data the covid-19 pandemic has affected many countries across the world, and disrupted the day to day activities of many people ...

Table 9: The contrastive synthetic data reweighted by MetaAdaptRank, where the top 3 cases are from Robust04 (News) and the last 3 cases come from TREC-COVID (BioMed). Their meta-weights are marked in the parenthesis ↑ (more important) and ↓ (down-weight). The **red** texts are the specific contents of the positive documents, and the **blue** texts are mentioned in both positive and negative documents. The document snippets are manually selected.

MetaAdaptRank with the hybrid source and regard the labeled data from previous rounds (R1-R4) as target data in meta-reweighting. Then we continuously train MetaAdaptRank using the labeled data from the previous rounds. In the training processes, we utilize Adam optimizer with a learning rate of $2e-5$. Both the batch size and the accumulation step are set to 8. In addition, to ensure a fair comparison with the submitted search systems, we post-process our results according to official guidelines.

A.3 Supplementary Case Studies

Table 9 shows more cases for the other two datasets, Robust04 (News) and TREC-COVID (BioMed), to verify the effectiveness of MetaAdaptRank in different domains. The first three cases are from Robust04, and the rest cases are from TREC-COVID.

For the first synthetic cases, our CTSyncSup can extract characteristic keywords, e.g., “radars” and “colombia”, from the positive documents to generate more informative queries, while SyncSup tends to capture general keywords to create broad queries, which may lack the ability to distinguish between different documents. Besides, CTSyncSup can extract some necessary themes from the specific documents, such as the particular time “1993” and the adjective “increased”, as shown in the second case.

Moreover, cases 4 and 5 show the effectiveness of our contrastive synthesis for biomedical domains. CTSyncSup can capture “lung” and “quarantine prevent” instead of general keywords, such as “sars cov” and “symptoms” often mentioned in COVID-related documents. These observations show that CTSyncSup can extract more specific information

to generate more informative and discriminative queries for different target domains.

We further explore those synthetic instances that are assigned zero weights by MetaAdaptRank, such as the third and sixth cases. In the third case, although CTSyncSup captures the two keywords “language” and “osvaldo rodriguez” from the positive document, its synthetic query is actually less relevant to the main topic of the positive document. For the sixth case, CTSyncSup fails to exclude the

phrase “covid-19 pandemic” related to both the positive and negative documents, which causes the synthetic query unable to distinguish between them. Fortunately, MetaAdaptRank can effectively identify the synthetic instances whose relevance matching patterns between synthetic queries and positive documents are unclear or non-unique and then precludes such misleading synthetic supervision data by assigning them zero weights.