

Task-level Distributionally Robust Optimization for Large Language Model-based Dense Retrieval

Guangyuan Ma^{1,2*}, Yongliang Ma³, Xing Wu^{1,2}, Zhenpeng Su^{1,2}, Ming Zhou³, Songlin Hu^{1,2†}

¹Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China

²School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China

³Langboat Technology, Beijing, China

{maguangyuan,wuxing,suzhenpeng,husonglin}@iie.ac.cn, {mayongliang,zhouming}@langboat.com

Abstract

Large Language Model-based Dense Retrieval (LLM-DR) optimizes over numerous heterogeneous fine-tuning collections from different domains. However, the discussion about its training data distribution is still minimal. Previous studies rely on empirically assigned dataset choices or sampling ratios, which inevitably lead to sub-optimal retrieval performances. In this paper, we propose a new task-level Distributionally Robust Optimization (tDRO) algorithm for LLM-DR fine-tuning, targeted at improving the universal domain generalization ability by end-to-end reweighting the data distribution of each task. The tDRO parameterizes the domain weights and updates them with scaled domain gradients. The optimized weights are then transferred to the LLM-DR fine-tuning to train more robust retrievers. Experiments show optimal improvements in large-scale retrieval benchmarks and reduce up to 30% dataset usage after applying our optimization algorithm with a series of different-sized LLM-DR models.

Code — <https://github.com/tdro-llm/tdro>

Datasets — <https://huggingface.co/tdro-llm>

Introduction

Dense retrieval (Karpukhin et al. 2020) recalls relevant documents from large-sized candidate pools with the similarity search (Mussmann and Ermon 2016) of query-passage embeddings. The recent bloom of Large Language Model-based Dense Retrieval (LLM-DR) (Wang et al. 2024a; Meng et al. 2024; Muennighoff et al. 2024) promotes remarkable retrieval abilities with better foundation models (Touvron et al. 2023; Bai et al. 2023; Jiang et al. 2023) and large-scale training collections (Wang et al. 2024a; Xiao et al. 2023).

LLM-DR fine-tuning learns LLM-based retrieval models with heterogeneous training collections (Reimers 2019) from multiple domains with different learning difficulties. The data distribution of training collections, i.e. a mixture of data with *chosen datasets* or *sampling ratio* on each dataset, significantly influences the general retrieval performances of dense retrievers (Oren et al. 2019; Meng et al. 2024). However, the choice or sampling ratio of the training sets still

*Work done during internship at Langboat Technology.

†Corresponding author.

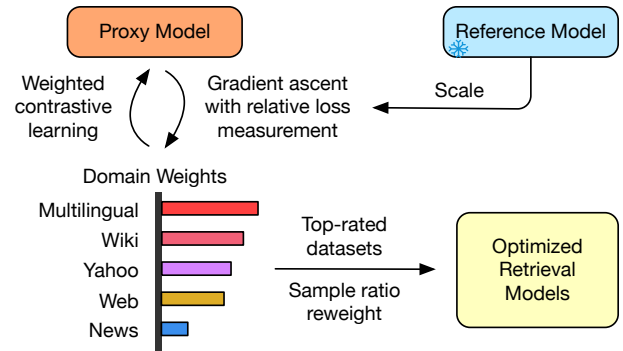


Figure 1: Task-level Distributionally Robust Optimization for Large Language Model-based Dense Retrieval.

relies heavily on intuitional assessments. It’s hard to decide whether an empirical assigned data sampling ratio is optimal for the models to perform well, i.e. robust to all tasks. Nowadays, the robustness of data distributional optimization for LLM-DR is still very limited.

Distributionally Robust Optimization (DRO) (Oren et al. 2019; Sagawa et al. 2019; Piratla, Netrapalli, and Sarawagi 2022) receives extensive discussions for battling unbalanced data composition. GroupDRO (Sagawa et al. 2019), the most popular algorithm for DRO, optimizes on the worst-case loss of the corresponding domain, which picks a domain with the highest loss at each training step and up-weights the loss of this domain. Although there was an attempt (Yu et al. 2022) to utilize vanilla GroupDRO or variants of DRO algorithm (Piratla, Netrapalli, and Sarawagi 2022) for dense retrieval fine-tuning, the optimization is limited to a small BERT-based model over clustered groups of one single dataset, i.e. MS-MARCO (Nguyen et al. 2016), which fails to generalize to LLM-DR fine-tuning with multiple heterogeneous training collections. It’s profitable to solve the data distribution issue of LLM-DR in an end-to-end manner like DRO, but such a study is still left for further exploration.

The existing DRO algorithms (Oren et al. 2019; Sagawa et al. 2019; Piratla, Netrapalli, and Sarawagi 2022), such as GroupDRO, are theoretically incompatible with LLM-DR fine-tuning due to *different batch sampling strategies* and *incommensurable loss scales*. Firstly, DRO requires all

Dataset	Loss
Yahoo answers (Title-Answer)	3.9257
MS-MARCO	1.3312
DuReader	0.6925

Table 1: Comparison of loss scales for Yahoo answers (Title-Answer) (Zhang, Zhao, and LeCun 2015), MS-MARCO (Nguyen et al. 2016), and DuReader (Qiu et al. 2022) datasets. The model used here is Qwen1.5-0.5B trained with uniform data sampling ratios for 1k steps.

domain data mandatory in one batch for end-to-end comparisons. It dynamically reweights the training losses based on the worst-case group and derives the robust-optimized proxy model as the final model directly. However, the LLM-DR fine-tuning collects the heterogeneous sets in a homogeneous batching method (Meng et al. 2024) during its fine-tuning, which means only one domain can be collected to the whole batch to ensure that in-batch negatives are sampled from the same task. What’s worse, the heterogeneous collections used by LLM-DRs have significantly different loss scales. If directly applying the DRO algorithms (Sagawa et al. 2019; Yu et al. 2022; Piratla, Netrapalli, and Sarawagi 2022) to LLM-DRs, the models will always bias towards the domain with the highest training loss, i.e. worst case loss, which will hurt the fine-tuning process. As is shown in Table 1, the loss of Yahoo answers with a Qwen1.5-0.5B retriever is three times over MS-MARCO and five times over DuReader. Directly using worst-case loss will make the model always biased towards Yahoo, rather than MS-MARCO or DuReader.

To tackle the above optimization problems for LLM-DR, we develop a new task-level Distributionally Robust Optimization (tDRO) algorithm that aims at improving the general domain adaption ability across multiple *tasks*¹. Firstly, to coordinate with different batch sampling strategies, our algorithm separates the DRO optimization and LLM-DR fine-tuning as is presented in Figure 1. Instead of directly learning a robust model (Oren et al. 2019; Sagawa et al. 2019), such separation first learns domain weights with a proxy model via the DRO algorithm and then transfers learned domain weights to LLM-DR fine-tuning. The proxy model is initialized from a small-sized LM, e.g. Qwen1.5-0.5B (Bai et al. 2023). It receives uniformly sampled training collections within an input batch, computes contrastive losses of each domain, and uses them as the gradients of domain weights. Then it transfers the learned weight or merely chooses the top-weighted datasets all LLM-DR model fine-tunings with different sizes. This separation shares several benefits: We can sample all domains within a batch in the tDRO stage and use task-homogeneous batching in the fine-tuning stage, which makes DRO work well while not hurting the final retrieval performances of LLM-DRs. Also, a small-sized LM can be used in tDRO optimizations to reduce com-

¹For simplicity, we treat each *dataset* as a single *domain / group / task* by following (Sagawa et al. 2019). We will mix the usage of these different expressions in the rest of our paper.

putational overheads.

Secondly, as discussed above, the heterogeneous domains have different loss scales. To make a comparable measurement of domain running losses, we use a trained LLM-DR model, e.g. Qwen1.5-0.5B with uniform data sampling ratios, as the reference model and forward it with the same inputs. We compute the *relative loss measurement* by dividing the proxy loss with reference loss. Intuitively, the relative loss measurement represents the improving headroom for the corresponding domains. Higher gradients will up-weight more corresponding domain weights.

To verify the effectiveness of the tDRO algorithm, we conduct data distribution optimization on open-sourced sentence transformers training data (Reimers 2019). We test on three large-scale retrieval benchmarks to fully assess the universal retrieval abilities across different languages and domains, including multilingual MIRACL (Zhang et al. 2023), cross-lingual MKQA (Longpre, Lu, and Daiber 2021), and monolingual BeIR (Thakur et al. 2021). Experiments shows steady improvements with less dataset usage after applying tDRO optimization.

Algorithm

Problem Statement

Assume the training collections D^{train} of LLM-DR fine-tuning are composed of k individual datasets, each of them is assigned with a domain weight α , representing a choice of probability distribution P_α over joint training collections:

$$P_\alpha = \sum_{g=1}^k \alpha_g U(D_g^{train}), s.t. \sum_{g=1}^k \alpha_g = 1. \quad (1)$$

U is a uniform distributional sampling function. And $\alpha_g U(D_g^{train})$ means sampling from such distribution with weight α_g for group g , which is also called α -covered probability distribution (Oren et al. 2019). The goal of LLM-DR data distributional optimization is to find an optimal distribution P_α , or a choice of weights α specifically, enabling the model to perform well on all downstream tasks D^{test} . Note that downstream tasks D^{test} are not necessarily identical to the fine-tuning sets D^{train} .

Task-level Distributionally Robust Optimization

The task-level Distributionally Robust Optimization (tDRO) algorithm parameterizes domain weights α and tries to learn the best choice of α for all tasks in an end-to-end manner. The tDRO holds a basic assumption for solving the data distributional optimization of LLM-DR: In a scenario with multiple independent and identically distributed (iid) task collections, if a model is robust to the training phase, then it will be robust to most of the test sets. Thus like most DRO algorithms (Oren et al. 2019; Sagawa et al. 2019; Piratla, Netrapalli, and Sarawagi 2022), tDRO operates on the robustness of training collections for universal adaption abilities. The whole optimization pipeline includes a separate tDRO stage and LLM-DR fine-tuning stage.

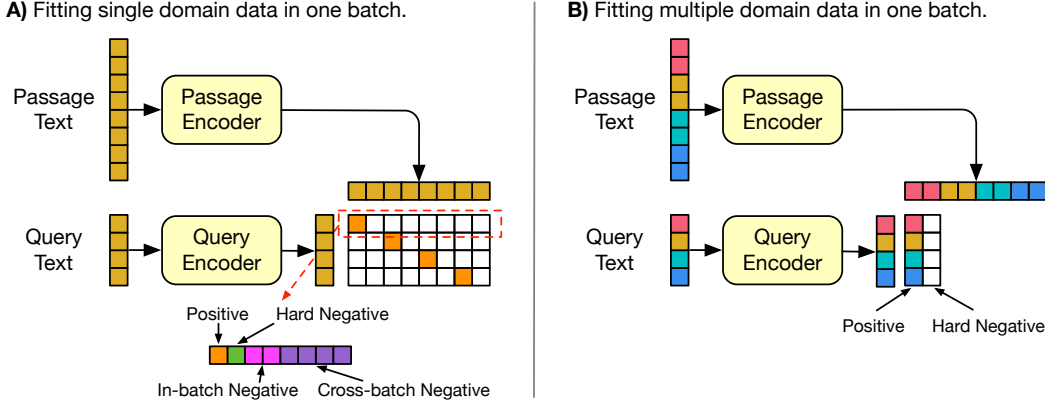


Figure 2: Different batch sampling strategies and negative types for A) LLM-DR Fine-tuning and B) Distributionally Robust Optimization.

InfoNCE Loss tDRO treats each dataset as a task (or domain) at minimal granularity and trains a proxy model with parameters θ_{proxy} to adaptively compare and update the domain weights at the task level. At each training step t , a batch of training data with domain sampling probabilities $1/k$ is forwarded through the proxy model. By extracting hidden states from the last token position, each batch item comprises a query representation q , a positive document representation d^+ , and several hard negative (HN) document representations d_{HN}^- . InfoNCE loss (van den Oord, Li, and Vinyals 2018), i.e. contrastive loss, is used to calculate losses of the proxy model for each group:

$$\mathcal{L}^{proxy} = -\log \frac{e^{q \cdot d^+ / \tau}}{e^{q \cdot d^+ / \tau} + \sum e^{q \cdot d_{HN}^- / \tau}}, \quad (2)$$

where τ is a contrastive temperature. In-batch negative sampling is not an option for tDRO because different tasks could induce false negatives and reduce negative qualities.

Optimization Objective tDRO learns domain weights α with a dual-optimization objective, which minimizes the supremum of the α -weighted sum of group loss measurements. Such an objective ensures universal robustness by lowering the upper bound of the worst groups:

$$\min_{\theta} \sup_{\alpha} \sum_{g=1}^k \alpha_g \mathcal{M}_g(\theta_{proxy}; q, d^+, d^-), \quad (3)$$

where *group loss measurement* \mathcal{M}_g is a scalar corresponding to averaged losses, representing the improving headrooms for each group. Following the previous DRO framework (Sagawa et al. 2019), the above object is optimized by interleaving the updates of weights α and proxy model parameters θ .

Weights Update For weights update, tDRO optimizes the above object by up-weighting the corresponding domain weights with exponential gradient ascent:

$$\alpha_g^t = \alpha_g^{t-1} e^{\eta_{\alpha} \mathcal{M}_g}, \quad (4)$$

where η_{α} is the learning rate for domain weights. As a common optimization practice, gradient normalization is used on the above gradients to ensure stable weight updates. Intuitively, a higher loss measurement induces more up-weighting of the corresponding group. After the update of α , a re-normalization is performed to ensure $\sum_{g=1}^k \alpha_g = 1$.

Relative Loss Measurement The key component of tDRO is how to derive a balanced and comparable loss measurement \mathcal{M}_g . GroupDRO directly uses the average group loss $\mathbb{E}(\mathcal{L}_g^{proxy})$ as the loss measurement, where \mathbb{E} is the arithmetic mean function. However, as presented in Table 1, the averaged contrastive losses of each group are not comparable. Directly using the average group loss will always make the proxy model biased toward one group with the highest loss. To solve the above issue, we introduce a trained reference model *ref*, forward it with the same inputs, compute reference loss as Formula (2), and *divide* the proxy loss with reference loss to dynamically scale the average group loss. This design is called *relative loss measurement*.

$$\mathcal{M}_g = \mathbb{E}(\mathcal{L}_g^{proxy}) / \mathbb{E}(\mathcal{L}_g^{ref}). \quad (5)$$

The reference model is frozen and will not be updated during the tDRO stage. In our implementation, we initialize the reference model with Qwen1.5-0.5B (Bai et al. 2023) and train with uniform sampling weights on all training sets. The training setting of the reference model follows the LLM-DR fine-tuning recipe, which will be described later.

Proxy Model Update After updating domain weights, the proxy model is updated with α -weighted contrastive loss.

$$\theta^t = \theta^{t-1} - \alpha_g \eta_{\theta} \nabla_{\theta}, \quad (6)$$

where η_{θ} is the learning rate of the proxy model, and ∇_{θ} is the gradient of proxy model obtained by backpropagation. The AdamW optimizer can be used for fast convergence.

LLM-DR Fine-tuning

LLM-DR fine-tuning also uses contrastive loss to pull positive representations together and push negative representa-

Dataset	Language	Category	Deduped Size
agnews	English	News	1,157,745
AllNLI	English	NLI	277,230
altlex	English	Wikipedia Pair	112,696
amazon_review_2018	English	Amazon	999,999
cnm_dailymail	English	News	311,971
codesearchnet	English	Github	1,375,067
dureader	Chinese	Web Collections	86,395
eli5_question_answer	English	Reddit	325,390
gooaq_pairs	English	Web Collections	3,012,347
hotpotqa	English	Wikipedia QA	85,000
medmcqa	English	Medical	160,865
miracl	16 languages	Multilingual Wikipedia	32,405
mr_tydi_combined	11 languages	Multilingual Wikipedia	48,475
msmarco	English	Web Collections	502,854
nq	English	Wikipedia QA	58,800
quora_duplicates_triplets	English	Forum Duplicates	97,011
searchQA_top5_snippets	English	Web Collections	117,219
sentence-compression	English	News	180,000
SimpleWiki	English	Wikipedia Pair	102,225
squad_pairs	English	Wikipedia QA	87,595
stackexchange_duplicates_title-body	English	Forum Duplicates	250,516
t2ranking	Chinese	Web Collections	200,376
trivia	English	Wikipedia QA	60,370
xsum	English	News	226,711
yahoo_answers_title_answer	English	Yahoo	1,198,018

Table 2: Training datasets information. Note that strict deduplication is performed on all training sets with Simhash (Manku, Jain, and Sarma 2007) to ensure no overlap between training and testing sets.

tions away. But it uses completely different batch sampling strategies and negative types from tDRO, which is one of the reasons that previous DRO algorithms like GroupDRO are theoretically incompatible with LLM-DR fine-tuning.

LLM-DR Batching Strategy As is shown in Figure 2A, LLM-DR fine-tuning fits the data from one single domain in each batch to ensure the quality of in-batch negatives. This batching strategy is also called task-homogeneous batching (Meng et al. 2024). As is a common practice, LLM-DR trains with large batch sizes, e.g. 2048 in our implementation, and three types of negatives, including hard negatives (HN), in-batch negatives (IBN), and cross-batch negatives (CBN). Large batch size enables contrastive learning of LLM-DR using more negatives, especially in-batch and cross-batch negatives. Hard negatives are provided by individual datasets, which are mined from existing retrieval systems like BM25 (Robertson et al. 1994) or dense retrievers (Karpukhin et al. 2020). In-batch negatives are samples from other data items within a batch. Cross-batch negatives are samples gathered from other GPU batches. Overall, the contrastive loss (\mathcal{L}^{CL}) for LLM-DR is formulated as follows.

$$\mathcal{L}^{CL} = -\log \frac{e^{q \cdot d^+ / \tau}}{e^{q \cdot d^+ / \tau} + \sum e^{q \cdot \{d_{HN}^-; d_{IBN}^-; d_{CBN}^- \} / \tau}}. \quad (7)$$

However, tDRO compares and updates domain weights in an end-to-end manner, thus it requires fitting all domain data in one batch. As displayed in Figure 2B, The training batch

is composed of all domains, while the in-batch/cross-batch negatives are not used in tDRO.

Experiments

Experiment Settings

Datasets A total of 25 individual datasets are used in our experiments, covering categories of Wikipedia, web collection, news, GitHub, yahoo, etc. Most of them are directly taken from sentence transformers training data (Reimers 2019). BGE-large-en-1.5 (Xiao et al. 2023) is used to mine negatives if the original datasets (several English datasets) have no negatives provided. Several popular multilingual datasets are also included in the training sets, including MIRACL (Zhang et al. 2023) and Mr.Tydi (Zhang et al. 2021). All information about datasets is listed in Table 2.

Settings For the tDRO algorithm, both the proxy and reference models are initialized from Qwen1.5-0.5B (Bai et al. 2023) for computational efficiency. tDRO is performed with a total batch size of 2048, query & document maximum sequence lengths of 128 & 512, a proxy model learning rate η_θ of 1e-4, contrastive temperature τ of 0.002, weights learning rate η_α of 2e-2, and seed of 42. The weights from the tDRO stage are directly transferred to LLM-DR fine-tuning. Two weight transfer strategies are utilized:

1. **Top-rated dataset selection:** Use the *Top-rated* datasets and discard the datasets with lower weights. This helps reduce dataset usage.

Benchmark (# Dataset)	MIRACL (18)		MKQA (25)		BeIR (15)
Metric	nDCG@10	Recall@100	Accuacy@10	Accuacy@100	nDCG@10
Uniform Sampling Baselines					
Qwen-0.5B	45.8	80.5	43.1	61.3	47.5
Qwen-1.8B	50.9	84.7	45.0	64.0	48.8
Qwen-4B	55.9	88.7	53.7	70.2	51.8
Qwen-7B	59.6	90.6	58.7	73.6	52.3
Mistral-7B	61.3	91.6	59.8	72.8	55.2
Llama3-8B	64.1	92.8	64.0	75.8	55.0
tDRO - Dataset Selection Top-70%					
Qwen-0.5B	48.7* (+2.9)	82.1* (+1.6)	45.4* (+2.3)	62.3* (+1.0)	48.9* (+1.4)
Qwen-1.8B	54.1* (+3.2)	86.6* (+1.9)	48.6* (+3.6)	65.6* (+1.6)	50.2* (+1.4)
Qwen-4B	58.6* (+2.7)	90.0* (+1.3)	57.0* (+3.3)	71.4* (+1.2)	52.6* (+0.8)
Qwen-7B	61.6* (+2.0)	91.4* (+0.8)	59.9* (+1.2)	73.8* (+0.2)	53.3* (+1.0)
Mistral-7B	63.8* (+2.5)	92.4* (+0.8)	62.5* (+2.7)	73.8* (+1.0)	55.2 (+0.0)
Llama3-8B	66.4* (+2.3)	93.5* (+0.7)	66.0* (+2.0)	76.4* (+0.6)	55.1 (+0.1)
Avg Gains	+2.6	+1.2	+2.5	+0.9	+0.8
tDRO - Sample Ratio Reweighting					
Qwen-0.5B	49.1* (+3.3)	82.7* (+2.2)	45.5* (+2.4)	62.2* (+0.9)	48.3* (+0.8)
Qwen-1.8B	53.6* (+2.7)	86.5* (+1.8)	50.5* (+5.5)	66.8* (+2.8)	49.7* (+0.9)
Qwen-4B	58.4* (+2.5)	90.0* (+1.3)	57.8* (+4.1)	72.0* (+1.8)	51.9 (+0.1)
Qwen-7B	61.0* (+1.4)	91.1 (+0.5)	59.8* (+1.1)	73.6 (+0.0)	52.4 (+0.1)
Mistral-7B	63.4* (+2.1)	92.4* (+0.8)	62.8* (+3.0)	74.0* (+1.2)	55.4 (+0.2)
Llama3-8B	66.3* (+2.2)	93.4* (+0.6)	67.0* (+3.0)	76.8* (+1.0)	55.0 (+0.0)
Avg Gains	+2.4	+1.2	+3.2	+1.3	+0.4

Table 3: Retrieval performances and corresponding gains of tDRO algorithm on MIRACL dev, MKQA test, and BeIR test benchmarks. The highest retrieval scores and average performance gains are marked as **bold**. *Significant improvements ($p \leq 0.01$) over the corresponding baseline. MS-MARCO in BeIR uses the dev split because there is no public test label.




Figure 3: Weights comparison among baseline, tDRO, and other loss measurement designs.

2. Sample Ratio Reweighting: Directly use the weights to *reweight* the sample ratios of datasets.

We conduct weights transfer on Qwen-1.5 0.5B, 1.8B, 4B, 7B (Bai et al. 2023), Mistral-0.1-7B (Jiang et al. 2023) and Llama3-8B (Touvron et al. 2023) for LLM-DR fine-tuning. Contrastive learning is performed with the same batch size, sequence lengths, model learning rate ($1e-4$), and contrastive temperature as stated before. Gradient cache (Gao et al. 2021), flash attention 2 (Dao 2023), full-shard data paral-

lel (FSDP), activation checkpointing and low-rank adapter (LoRA) (Hu et al. 2022) with $r = 8$, $\alpha = 32$ and dropout ratio of 0.1 are used to reduce GPU memory usage. Following the previous work (Wang et al. 2024a; Su et al. 2023), prompt instructions are added on the query side for multi-tasks during training and evaluation. All trainings are performed on 8 NVIDIA H800 GPUs with 4.5 hours for tDRO and less than 10 hours for all LLM-DR fine-tunings.

Benchmark	Enhance	MIRACL	MKQA	BeIR
Metric → Models ↓	Special Pre-train	nDCG@10	Acc@100	nDCG@10
BM25		31.9	39.9	41.7
mContriever	✓	43.1	67.9	46.0
mE5-large-inst	✓	64.6	70.2	52.3
E5-Mistral*		62.2	68.6	52.7
tDRO + LLM-DR Llama3-8B				
Top-70%		66.4	76.4	55.1
Reweighting		66.3	76.8	55.0

Table 4: Retrieval scores among state-of-the-art retrievers. *We take the released BeIR scores without GPT data for fair comparisons.

Baselines and Benchmarks To compare the performance changes after tDRO optimization, we choose LLM-DR with *uniform sampling weights*² as baselines. All other settings are kept the same. Note that results from some recent state-of-the-art retrievers, including BM25 (Robertson et al. 1994), mContriever/Contriever (Izacard et al. 2021), E5-Mistral-7b (Wang et al. 2024a), and Multilingual-e5-large-instruct (Wang et al. 2024b) are also added to our results. But we DO NOT seek to compare with them. Some of them utilize multiple training enhancements, such as data synthesis with ChatGPT and special pre-trainings (Wang et al. 2023; Wu et al. 2023; Ma et al. 2024a) for retrieval, which is an unfair comparison to our experiments and out of the research scope for our paper.

MIRACL (Zhang et al. 2023), MKQA (Longpre, Lu, and Daiber 2021), and BeIR (Thakur et al. 2021) are used as main retrieval benchmarks. MIRACL is a huge multilingual retrieval benchmark across 18 languages with 13K queries and 90M passages. Cross-lingual benchmark MKQA holds 25 different languages of parallel Wikipedia queries. Following (Chen et al. 2024), the cross-lingual retrieval is performed by using queries (6.6k for each) of different languages to recall relevant English Wikipedia with 2.7M passages³. BeIR is a heterogeneous English retrieval benchmark with 15 public datasets and 34M documents, covering a wide range of retrieval domains and text symmetries. MIRACL is reported with nDCG@10 and Recall@100. MKQA is reported with Accuracy@{10, 100} (Acc). Examining performances at threshold 10 assess the top-ranking ability and threshold 100 for recalling capacity at a larger window. BeIR is reported with nDCG@10 by following the original paper (Thakur et al. 2021). The whole evaluation takes around 30 hours with 8 NVIDIA H800 GPUs for the largest LLM-DR retriever, Llama3-8B.

Results

Retrieval benchmarks on multilingual, cross-lingual, and monolingual English retrieval are conducted and listed in

²A total of 25 datasets are used in our experiments, thus the sampling weight for each dataset is $1/25 = 0.04$

³<https://huggingface.co/datasets/Tevatron/wikipedia-nq-corpus>

	BEIR	MIRACL	MKQA
Metric	nDCG@10	nDCG@10	Acc@100
Qwen-0.5B-Baseline	47.5	45.8	61.3
Sampling Strategy Choices for tDRO			
<i>Dataset Selection</i>			
w/ Bottom-50%	45.2 _{-2.3}	44.6 _{-1.2}	60.7 _{-0.6}
w/ Top-60%	48.4 _{+0.9}	48.2 _{+2.4}	62.0 _{+0.7}
w/ Top-70%	48.9 _{+1.4}	48.7 _{+2.9}	62.3 _{+1.0}
w/ Top-80%	48.9 _{+1.4}	47.8 _{+2.0}	62.1 _{+0.8}
<i>Sampling Ratio Reweighting</i>	48.3 _{+0.8}	49.1 _{+3.3}	62.2 _{+0.9}
<i>Loss Reweighting</i>	48.2 _{+0.7}	46.3 _{+0.5}	61.2 _{-0.1}
Loss Measurements (w/ Sampling Ratio Reweighting)			
$\bar{\mathcal{L}}^{proxy} / \bar{\mathcal{L}}^{ref}$	48.3 _{+0.8}	49.1 _{+3.3}	62.2 _{+0.9}
$\bar{\mathcal{L}}^{proxy} - \bar{\mathcal{L}}^{ref}$	47.5 _{-0.0}	45.3 _{-0.5}	60.1 _{+1.2}
$\bar{\mathcal{L}}^{proxy}$	47.2 _{-0.3}	47.0 _{+1.2}	61.5 _{+0.2}

Table 5: Ablation studies on sampling strategy choices and group loss designs.

Table 3. tDRO optimized domain weights are listed in Figure 3. Top-70% is the optimal line for the top-rated dataset selection strategy. Ablation on the percentages of dataset selection will be presented later.

tDRO Boosts Retrieval Performances. tDRO boosts the universal retrieval performances of LLM-DR on a series of different-sized base LLM, including Qwen-1.5 (0.5B, 1.8B, 4B, 7B), Mistral-0.1 7B, and Llama3 8B. Both the top-rated dataset selection and the sample ratio reweighting strategies work well on most retrieval benchmarks with significant performance gains. The improvements in multilingual MIRACL and cross-lingual MKQA retrieval are attributed to the up-weighting of multilingual datasets, e.g. MIRACL.

tDRO Balances Data Usages As is shown in Figure 3, although multilingual training sets are less than monolingual sets, tDRO balances the lingual distribution differences by top-rating the multilingual MIRACL and T2-Ranking. In contrast, tDRO down-weights 9 monolingual English datasets (Nearly 30% amounts) to less than 0.01. If we remove less-weighted datasets and keep only the Top-70% sets, the results on multilingual, cross-lingual, and monolingual retrieval benchmarks even get significantly better, although less data is involved in LLM-DR fine-tuning. This is attributed to the end-to-end task-level data distribution optimization by tDRO, which balances the data distribution and helps LLM-DR training to focus on more useful and challenging data. Notably, as shown in Table 4, tDRO-optimized LLM-DR gets leading performances without using further enhancements, such as GPT4 data synthesis and special pre-training.

Analysis

Weight Transferring Strategies

According to Table 5, top-rated dataset selection reaches peak performances with Top-70% datasets. On the contrary, using Bottom-50% significantly hurts the performances. This strategy is stable on nearly all benchmarks. However, as is shown in Table 3, some monolingual results seem to reach peak levels without further gains, such as Mistral-7B and Llama3-8B for BeIR. This phenomenon only occurs in LLMs with larger parameters ($>7B$) on BeIR. A potential explanation could be that larger LLMs have stronger capacities to tolerate the data weight changes when plenty of monolingual data is provided. Under this circumstance, large LLMs can perform evenly well.

Sample ratio reweighting has more performance gains on MKQA than the selection strategy because it incurs more sampling probabilities on multilingual datasets, e.g. MIRACL and T2-Ranking, as displayed in Figure 3. However, such over-weighting on multilingual datasets is also the reason for no significant gains on monolingual BeIR of large LLMs, e.g. Qwen-4B, 7B, Mistral-7B, and Llama3-8B.

Additionally, we also test re-scaling contrastive losses of LLM-DR with transferred weights. However, such loss reweighting improves BeIR but has no gains on MIRACL and MKQA. A potential reason could be that loss reweighting only changes the loss scale, but does not import more multilingual data because the data sampling ratios are unchanged.

Loss Measurement Designs

Loss measurement is the key design to get proper domain loss scales. tDRO utilizes *relative loss measurement* by dividing the proxy loss with the reference loss. We also tested two additional choices:

1. Minus the proxy loss with reference loss. This design is also called group excess loss in DRO (Oren et al. 2019). But obviously, the minus operation could not *scale* the big loss differences of heterogeneous collections for LLM-DR fine-tuning.
2. Directly using proxy loss.

As is shown in Table 5, these two loss measurements hurt the performances. This is attributed to the incomparable loss scales of different datasets. According to Figure 3, both measurements over-prefer Yahoo significantly, because of the biggest loss scale of Yahoo as shown in Table 1.

Related Works

Large Language Model-based Dense Retrieval (LLM-DR)

Dense retrieval (DR) has gained significant improvements in recent years because of the rapid developments of large language models (LLM) (Ouyang et al. 2022; Touvron et al. 2023; Bai et al. 2023; Jiang et al. 2023) and growing collections of heterogeneous training data (Reimers 2019). In short, the model parameter is growing, and the datasets are increasing. Sentence-T5 (Ni et al. 2022) first scales the

foundation models to billion parameter T5 and gains good abilities of sentence embeddings. RepLLaMA and RankLLaMA (Ma et al. 2024b) first finetune retrievers and re-rankers on Llama2 with a single MS-MARCO dataset and get remarkable improvements over a series of small-sized baseline retrievers. E5-Mistral (Wang et al. 2024a) utilizes 500k GPT3.5/4 synthesized training data and 15 well-collected datasets to fine-tune the retriever, further pushing the boundary of retrieval abilities. It also transfers these data to specially pre-trained small-sized model mE5 (Wang et al. 2024b) for state-of-the-art (SOTA) multilingual performances. However, the utilization of training collection for LLM-DR still relies on intuitional assessments. As discussed in this paper, empirical assigned data choices or sampling ratios, e.g. uniform sampling, incur sub-optimal performances.

Distributionally Robust Optimization (DRO)

Distributionally Robust Optimization (DRO) is an effective way to battle unbalanced data distributions. Topic CVaR (Oren et al. 2019), first proposes to minimize the worst-case loss of each topic. (Sagawa et al. 2019) proposes a gradient-based GroupDRO algorithm and successfully improves the worse-group accuracies. DoReMi (Xie et al. 2023) utilizes GroupDRO in LLM pre-training with minimax optimization for improving perplexities and downstream accuracies. CGD algorithm (Piratla, Netrapalli, and Sarawagi 2022) introduces the inter-group interactions into GroupDRO by substituting the worst-case loss with the inner product score of group gradients.

Research on DRO for dense retrieval is still minimal. COCO-DR (Yu et al. 2022) combines pre-training method coCondenser (Gao and Callan 2022) and CGD algorithm to battle the distribution shifts on both the document and query side. However, its scope is limited to small-sized BERT models over clustered groups of one single MS-MARCO (Nguyen et al. 2016) dataset. Our study aims to solve the optimization problem of data distribution in LLM-DR fine-tuning, which end-to-end reweights each dataset for optimal performance.

Conclusion

Large language model-based dense retrieval (LLM-DR) utilizes multiple heterogeneous training datasets in fine-tuning. Previous studies rely on empirical assessments to decide the sampling ratios of each dataset, which incurs unbalanced training data distribution and leads to sub-optimal performances. In our study, we propose a new task-level Distributionally Robust Optimization (tDRO) algorithm to improve the domain generalization ability by end-to-end reweighting the data distribution of each dataset. Experiments on large-scale retrieval benchmarks show steady improvements with less dataset usage.

References

Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; Hui, B.; Ji, L.; Li, M.; Lin, J.; Lin, R.; Liu, D.; Liu, G.; Lu, C.; Lu, K.; Ma, J.; Men,

- R.; Ren, X.; Ren, X.; Tan, C.; Tan, S.; Tu, J.; Wang, P.; Wang, S.; Wang, W.; Wu, S.; Xu, B.; Xu, J.; Yang, A.; Yang, H.; Yang, J.; Yang, S.; Yao, Y.; Yu, B.; Yuan, H.; Yuan, Z.; Zhang, J.; Zhang, X.; Zhang, Y.; Zhang, Z.; Zhou, C.; Zhou, J.; Zhou, X.; and Zhu, T. 2023. Qwen Technical Report. *CoRR*, abs/2309.16609.
- Chen, J.; Xiao, S.; Zhang, P.; Luo, K.; Lian, D.; and Liu, Z. 2024. BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation. *CoRR*, abs/2402.03216.
- Dao, T. 2023. FlashAttention-2: Faster Attention with Better Parallelism and Work Partitioning. *CoRR*, abs/2307.08691.
- Gao, L.; and Callan, J. 2022. Unsupervised Corpus Aware Language Model Pre-training for Dense Passage Retrieval. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2843–2853. Dublin, Ireland: Association for Computational Linguistics.
- Gao, L.; Zhang, Y.; Han, J.; and Callan, J. 2021. Scaling Deep Contrastive Learning Batch Size under Memory Limited Setup. In Rogers, A.; Calixto, I.; Vulic, I.; Saphra, N.; Kassner, N.; Camburu, O.; Bansal, T.; and Shwartz, V., eds., *Proceedings of the 6th Workshop on Representation Learning for NLP, RePLNLP@ACL-IJCNLP 2021, Online, August 6, 2021*, 316–321. Association for Computational Linguistics.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Izacard, G.; Caron, M.; Hosseini, L.; Riedel, S.; Bojanowski, P.; Joulin, A.; and Grave, E. 2021. Towards Unsupervised Dense Information Retrieval with Contrastive Learning. *CoRR*, abs/2112.09118.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; de Las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023. Mistral 7B. *CoRR*, abs/2310.06825.
- Karpukhin, V.; Oguz, B.; Min, S.; Lewis, P.; Wu, L.; Edunov, S.; Chen, D.; and Yih, W.-t. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 6769–6781. Online: Association for Computational Linguistics.
- Longpre, S.; Lu, Y.; and Daiber, J. 2021. MKQA: A Linguistically Diverse Benchmark for Multilingual Open Domain Question Answering. *Trans. Assoc. Comput. Linguistics*, 9: 1389–1406.
- Ma, G.; Wu, X.; Lin, Z.; and Hu, S. 2024a. Drop your Decoder: Pre-training with Bag-of-Word Prediction for Dense Passage Retrieval. In Yang, G. H.; Wang, H.; Han, S.; Hauff, C.; Zuccon, G.; and Zhang, Y., eds., *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, 1818–1827. ACM.
- Ma, X.; Wang, L.; Yang, N.; Wei, F.; and Lin, J. 2024b. Fine-Tuning LLaMA for Multi-Stage Text Retrieval. In Yang, G. H.; Wang, H.; Han, S.; Hauff, C.; Zuccon, G.; and Zhang, Y., eds., *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, 2421–2425. ACM.
- Manku, G. S.; Jain, A.; and Sarma, A. D. 2007. Detecting near-duplicates for web crawling. In Williamson, C. L.; Zurko, M. E.; Patel-Schneider, P. F.; and Shenoy, P. J., eds., *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, 141–150. ACM.
- Meng, R.; Liu, Y.; Joty, S. R.; Xiong, C.; Zhou, Y.; and Yavuz, S. 2024. SFR-Embedding-Mistral: Enhance Text Retrieval with Transfer Learning.
- Muennighoff, N.; Su, H.; Wang, L.; Yang, N.; Wei, F.; Yu, T.; Singh, A.; and Kiela, D. 2024. Generative Representational Instruction Tuning. *CoRR*, abs/2402.09906.
- Mussmann, S.; and Ermon, S. 2016. Learning and Inference via Maximum Inner Product Search. In Balcan, M.; and Weinberger, K. Q., eds., *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, 2587–2596. JMLR.org.
- Nguyen, T.; Rosenberg, M.; Song, X.; Gao, J.; Tiwary, S.; Majumder, R.; and Deng, L. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In Besold, T. R.; Bordes, A.; d’Avila Garcez, A. S.; and Wayne, G., eds., *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016*, volume 1773 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Ni, J.; Ábrego, G. H.; Constant, N.; Ma, J.; Hall, K. B.; Cer, D.; and Yang, Y. 2022. Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, 1864–1874. Association for Computational Linguistics.
- Oren, Y.; Sagawa, S.; Hashimoto, T. B.; and Liang, P. 2019. Distributionally Robust Language Modeling. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 4226–4236. Association for Computational Linguistics.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.

- Piratla, V.; Netrapalli, P.; and Sarawagi, S. 2022. Focus on the Common Good: Group Distributional Robustness Follows. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Qiu, Y.; Li, H.; Qu, Y.; Chen, Y.; She, Q.; Liu, J.; Wu, H.; and Wang, H. 2022. DuReader-Retrieval: A Large-scale Chinese Benchmark for Passage Retrieval from Web Search Engine. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 5326–5338. Association for Computational Linguistics.
- Reimers, N. 2019. Sentence Transformers Embedding Training Data. <https://huggingface.co/datasets/sentence-transformers/embedding-training-data>.
- Robertson, S. E.; Walker, S.; Jones, S.; Hancock-Beaulieu, M.; and Gatford, M. 1994. Okapi at TREC-3. In Harman, D. K., ed., *Proceedings of The Third Text REtrieval Conference, TREC 1994, Gaithersburg, Maryland, USA, November 2-4, 1994*, volume 500-225 of *NIST Special Publication*, 109–126. National Institute of Standards and Technology (NIST).
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2019. Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization. *CoRR*, abs/1911.08731.
- Su, H.; Shi, W.; Kasai, J.; Wang, Y.; Hu, Y.; Ostendorf, M.; Yih, W.; Smith, N. A.; Zettlemoyer, L.; and Yu, T. 2023. One Embedder, Any Task: Instruction-Finetuned Text Embeddings. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, 1102–1121. Association for Computational Linguistics.
- Thakur, N.; Reimers, N.; Rücklé, A.; Srivastava, A.; and Gurevych, I. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. *CoRR*, abs/2104.08663.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *CoRR*, abs/2302.13971.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR*, abs/1807.03748.
- Wang, L.; Yang, N.; Huang, X.; Jiao, B.; Yang, L.; Jiang, D.; Majumder, R.; and Wei, F. 2023. SimLM: Pre-training with Representation Bottleneck for Dense Passage Retrieval. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 2244–2258. Association for Computational Linguistics.
- Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2024a. Improving Text Embeddings with Large Language Models. *CoRR*, abs/2401.00368.
- Wang, L.; Yang, N.; Huang, X.; Yang, L.; Majumder, R.; and Wei, F. 2024b. Multilingual E5 Text Embeddings: A Technical Report. *CoRR*, abs/2402.05672.
- Wu, X.; Ma, G.; Lin, M.; Lin, Z.; Wang, Z.; and Hu, S. 2023. ConTextual Masked Auto-Encoder for Dense Passage Retrieval. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 4738–4746. AAAI Press.
- Xiao, S.; Liu, Z.; Zhang, P.; and Muennighof, N. 2023. C-Pack: Packaged Resources To Advance General Chinese Embedding. *CoRR*, abs/2309.07597.
- Xie, S. M.; Pham, H.; Dong, X.; Du, N.; Liu, H.; Lu, Y.; Liang, P.; Le, Q. V.; Ma, T.; and Yu, A. W. 2023. DoReMi: Optimizing Data Mixtures Speeds Up Language Model Pre-training. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Yu, Y.; Xiong, C.; Sun, S.; Zhang, C.; and Overwijk, A. 2022. COCO-DR: Combating Distribution Shifts in Zero-Shot Dense Retrieval with Contrastive and Distributionally Robust Learning. *CoRR*, abs/2210.15212.
- Zhang, X.; Ma, X.; Shi, P.; and Lin, J. 2021. Mr. TyDi: A Multi-lingual Benchmark for Dense Retrieval. In Ataman, D.; Birch, A.; Conneau, A.; Firat, O.; Ruder, S.; and Sahin, G. G., eds., *Proceedings of the 1st Workshop on Multilingual Representation Learning*, 127–137. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Zhang, X.; Thakur, N.; Ogundepo, O.; Kamaloo, E.; Alfonso-Hermelo, D.; Li, X.; Liu, Q.; Rezagholizadeh, M.; and Lin, J. 2023. MIRACL: A Multilingual Retrieval Dataset Covering 18 Diverse Languages. *Trans. Assoc. Comput. Linguistics*, 11: 1114–1131.
- Zhang, X.; Zhao, J. J.; and LeCun, Y. 2015. Character-level Convolutional Networks for Text Classification. In Cortes, C.; Lawrence, N. D.; Lee, D. D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, 649–657.