

Influence Guided Sampling for Domain Adaptation of Text Retrievers

Meet Doshi*, Vishwajeet Kumar, Yulong Li, Jaydeep Sen
IBM Research

Abstract

General-purpose open-domain dense retrieval systems are usually trained with a large, eclectic mix of corpora and search tasks. How should these diverse corpora and tasks be sampled for training? Conventional approaches sample them uniformly, proportional to their instance population sizes, or depend on human-level expert supervision. It is well known that the training data sampling strategy can greatly impact model performance. However, how to find the optimal strategy has not been adequately studied in the context of embedding models. We propose **Inf-DDS**, a novel reinforcement learning–driven sampling framework that adaptively reweighs training datasets guided by influence-based reward signals and is much lighter w.r.t. GPU consumption. Our technique iteratively refines the sampling policy, prioritizing datasets that maximize model performance on a target development set. We evaluate the efficacy of our sampling strategy in a wide range of text retrieval tasks, demonstrating strong improvements in retrieval performance and better adaptation compared to existing gradient-based sampling methods, while also being $1.5\times - 4\times$ cheaper in GPU compute. Our sampling strategy achieves an improvement in **5.03** absolute NDCG@10 while training a multilingual model *bge-m3* and an improvement in **0.94** absolute NDCG@10 while training *all-MiniLM-L6-v2*, even when starting from expert-assigned weights on a large pool of training datasets.

1 Introduction

Text-to-embedding dense retrievers have recently gained popularity due to their strong performance across diverse benchmarks (Karpukhin et al., 2020; Izacard et al., 2022; Wang et al., 2023; Chen et al., 2024; BehnamGhader et al., 2024). Their cross-domain generalization ability is typically achieved




Figure 1: Trade-off between training time (GPU hours) and average NDCG@10 on the MLDR-13 test set when optimizing *bge-m3-dense* on the MLDR-13 dev set.

through training on large heterogeneous datasets. For example, the *Sentence-Transformers* project trains generic embeddings using billions of instances drawn from domains with highly imbalanced data sizes. However, increasing dataset size alone does not necessarily improve embedding quality, making dataset selection and sampling strategy critical for learning robust representations.

While random sampling is a common default, it ignores differences in dataset informativeness when training on large data mixtures. Alternatives such as temperature sampling, proportional sampling, and manually tuned expert weights attempt to address this issue, but remain static and often require costly iterative experimentation. Since the optimal sampling distribution is generally unknown, dynamic adaptive sampling may better identify informative datasets during training.

There has consequently been substantial interest in adaptive sampling. Gradient-based approaches such as DDS (Wang et al., 2020a) and its multi-target extension (Wang et al., 2020b) use gradient-derived rewards to adjust the training distribution online. DoGE (Fan et al., 2024) proposes a *generalization estimation function* to approximate data influence, while methods like DoReMi (Xie et al.,

*Work done while at IBM Research. Correspondence: meetdoshi90@gmail.com.

2023a; Engstrom et al., 2024) rely on proxy models to estimate dataset utility. In practice, these dynamic methods face two main challenges: (1) instability and high variance introduced by stochastic gradients, which we empirically demonstrate in Section 5, and (2) substantial computational overhead from proxy models or expensive estimators. Together, these limitations motivate an online adaptive optimization strategy that learns sampling weights efficiently and robustly while remaining computationally tractable.

In this paper, we propose **Influence-guided Dynamic Data Sampling** strategy (**Inf-DDS**), a computationally efficient algorithm for domain-adaptive data sampling that overcomes limitations of existing gradient-based methods. Inf-DDS iteratively takes small gradient updates on each domain’s data while monitoring the impact on the downstream metric. Domains yielding larger performance improvements receive higher rewards, adaptively focusing training on the most informative data subsets. Inspired by recent influence-based methods (Koh and Liang, 2017; Bae et al., 2022; Fan et al., 2024; Yu et al., 2024; Xia et al., 2024), our method offers three key advantages: (1) it avoids noisy gradient estimates for reward computation, (2) it reuses computations from updating the sampling distribution parameters ψ also to update model parameters θ , making training substantially more efficient, and (3) it produces more reliable sampling trajectories for improved downstream performance.

Our contributions in this work are as follows: **a)** We propose Inf-DDS, an influence-guided reinforcement learning approach that learns to adjust sampling probabilities across diverse training datasets, improving target-domain retrieval performance while being much more computationally efficient. **b)** We validate Inf-DDS against robust benchmarks, including BEIR datasets (Thakur et al., 2021), Sentence-Transformers *all-MiniLM-L6-v2* (Reimers and Gurevych, 2019) and MLDR (Chen et al., 2024), demonstrating significant improvements while optimizing for a target domain. **Code:** github.com/meetdoshi90/InfDDS.

2 Related Work

Recent work on training models with large and diverse data pools has explored both heuristic sampling and adaptive reweighting schemes. A common approach samples domains uniformly or

via temperature-scaled distributions interpolating between uniform and size-proportional sampling. In multilingual pretraining, Cooldown (Li et al., 2024b) improves performance by oversampling high-resource languages early in training before gradually shifting toward uniform sampling to improve generalization. DoReMi optimizes domain sampling weights using the loss gap between proxy and target models.

Earlier domain adaptation methods identify relevant subsets of data using cross-entropy differences or lightweight classifiers (Moore and Lewis, 2010; Brown et al., 2020). DSIR (Xie et al., 2023b) performs importance sampling by assigning weights to instances based on hashed feature representations, while CRISP (Grangier et al., 2025b) clusters training data and assigns importance weights according to the frequency of source and target examples within clusters.

Recent influence estimation approaches prioritize examples that most affect model predictions or downstream performance (Grosse et al., 2023; Nikdan et al., 2025; Zhou et al., 2025). Methods such as LESS (Xia et al., 2024) and Quad (Zhang et al., 2025) approximate influence using first-order Taylor expansions of the target loss change after a gradient update. Related work further shows that data models can predict the influence of training samples on held-out targets (Ilyas et al., 2022; Engstrom et al., 2024). MATES (Yu et al., 2024) similarly estimates oracle influence using a small proxy model after a single update step, but small proxies may provide limited accuracy. In contrast, we leverage online proxy models to compute exact influence scores for learning sampling weights during training.

DDS and DoGE extend beyond fixed heuristics through bi-level optimization, jointly learning scorer networks and model parameters to prioritize examples whose gradients align with a held-out development set. MultiDDS (Wang et al., 2020b) further extends DDS to multiple targets through per-language scoring functions. However, gradient-based rewards often exhibit high variance in practice. Our method instead replaces noisy gradient-based rewards with online-computed influence scores derived from updated model parameters, yielding a simpler and more stable approach while reusing intermediate computations for efficiency. We refer interested readers to (Grangier et al., 2025a; Choi et al., 2023; Chung et al., 2023; Cao et al., 2024; Pruthi et al., 2020; Park et al.,

2023; Liu et al., 2025; Grosse et al., 2023) for broader discussions on domain adaptation and universal generalization.

3 Influence Guided Dynamic Data Sampling

In this section, we propose a reinforcement learning-based strategy for domain adaptation of text retrievers. We frame dataset sampling as a bilevel optimization problem and learn an adaptive policy to maximize performance on target datasets. Our approach is illustrated in Fig. 2 and elaborated in Algorithm 1.

3.1 Problem Formulation

Efficient data sampling for training text retrieval models: Our goal is to devise an adaptive sampling policy that efficiently samples batches from a large pool of training domains/datasets (M) to maximize the model performance on a set of target datasets (N) a.k.a. development or dev sets. We treat each training and dev dataset as a single homogeneous unit and optimize sampling at the dataset level.

Given a set of training datasets $\{\mathcal{D}_{\text{train}}^i\}_{i=1}^M$ with initial sampling probability as $P_D(i)$; $i \in 1, \dots, M$, a set of dev datasets $\{\mathcal{D}_{\text{dev}}^j\}_{j=1}^N$ on which we want our model to adapt, and initial model with parameters θ , our objective is to optimize the model parameter θ by learning a dynamic sampling strategy for P_D using a parameterized policy ψ that maximizes performance on development sets.

We formalize our objective through the following optimization problem:

$$\theta^*, \psi^* = \arg \min_{\theta, \psi} \mathbb{E}_{i \sim P_D(i; \psi)} [J(\theta, \mathcal{D}_{\text{train}}^i)] \quad (1)$$

where $J(\theta, \mathcal{D}_{\text{train}}^i)$ is the empirical risk.

The training datasets sampling probability distribution $P_D(i; \psi)$ is computed as:

$$P_D(i; \psi) = \frac{e^{\psi_i}}{\sum_{k=1}^M e^{\psi_k}} \quad (2)$$

where M is the number of datasets in the pool. This formulation allows the sampling strategy to dynamically adjust the probabilities based on the learned parameters (ψ), guiding the selection of datasets in a way that optimizes the model performance on the dev set. (ψ_i) represents the importance score associated with the train dataset i .

We assume that the dev sets $\mathcal{D}_{\text{dev}}$ have a distribution similar to the test set $\mathcal{D}_{\text{test}}$ that is, $\mathcal{D}_{\text{dev}} \approx \mathcal{D}_{\text{test}}$. Assuming the existence of N development datasets, Eq. (1) can be expanded as follows:

$$\begin{aligned} \psi^* &= \arg \min_{\psi} \frac{1}{N} \sum_{j=1}^N J(\theta^*(\psi), \mathcal{D}_{\text{dev}}^j) \\ \theta^* &= \arg \min_{\theta} \mathbb{E}_{i \sim P_D(i; \psi)} [J(\theta, \mathcal{D}_{\text{train}}^i)] \end{aligned} \quad (3)$$

This formulation naturally leads to a bi-level optimization problem involving θ and ψ , which can be effectively addressed by alternating optimization steps. Specifically, the model parameters θ are updated using the standard gradient descent algorithm, while the scorer parameters ψ , are optimized using the REINFORCE algorithm (Williams, 1992).

3.2 Influence based Rewards

Intuitively, the reward signal should guide the parameterized scoring network (ψ) to up-sample training datasets most likely to improve model performance on the development sets. We therefore propose an influence-based reward mechanism that quantifies the contribution of each training dataset to performance on a held-out development set. Existing influence-based methods (Xia et al., 2024; Fan et al., 2024; Yu et al., 2024) either rely on first-order approximations or simulate this contribution using smaller proxy models. In contrast, we employ online proxy models to obtain accurate influence scores. Specifically, at each time step t , we update the model parameters θ_t using dataset $\mathcal{D}_{\text{train}}^i$ for l gradient steps, where l is the minimum number of steps required to reach a meaningful local minima that demonstrates the potential benefit of up-sampling $\mathcal{D}_{\text{train}}^i$.

During each of these l steps, we estimate the influence of every training subset $\mathcal{D}_{\text{train}}^i$ on the current model θ_t . We do this by (1) taking l gradient steps on $\mathcal{D}_{\text{train}}^i$ to produce θ_{t+1}^i , (2) evaluating both θ_t and θ_{t+1}^i on each dev batch using an influence metric $\mathcal{M}$, and (3) computing the change in performance:

$$\Delta \mathcal{M}_j^i = \mathcal{M}(\theta_{t+1}^i; \mathcal{D}_{\text{val}}^j) - \mathcal{M}(\theta_t; \mathcal{D}_{\text{val}}^j) \quad (4)$$

where $\Delta \mathcal{M}_j^i \propto \mathcal{I}_{\theta}(i; \mathcal{D}_{\text{val}})$, and $\mathcal{I}_{\theta}(i; \mathcal{D}_{\text{val}})$ is the influence estimate that measures the effect of up-weighting $\mathcal{D}_{\text{train}}^i$ on $\mathcal{D}_{\text{val}}$.

To ensure robustness and stability in our estimates, we normalize across the development datasets by taking a mean over all influences $\bar{\Delta \mathcal{M}}^i$ where $\bar{\Delta \mathcal{M}}^i = \frac{1}{N} \sum_{j=1}^N \Delta \mathcal{M}_j^i$. We iterate the




Figure 2: Overview of Inf-DDS. The trainable scorer ψ and model parameters θ are optimized by generating online proxy models to compute influence scores, which serve as rewards for updating the scorer ψ . Proxy model gradients are efficiently reused for a weighted Reptile update on θ .

process l times over all train-dev pairs (i, j) , accumulating the reward values $\Delta \mathcal{M}^i$ to compute the final influence $\mathcal{I}^i$, which serves as a reliable measure of the impact of the i^{th} training dataset on the model’s performance on the development set.

Algorithm 1 Inf-DDS

```

1: Input:  $\{\mathcal{D}_{\text{train}}^i\}_{i=1}^M, \{\mathcal{D}_{\text{val}}^j\}_{j=1}^N$ , influence metric  $\mathcal{M}$ , inner steps per meta-update  $k$ 
2: Output: converged model  $\theta^*$ 
3: Initialize  $P_D(i; \psi, \tau) \leftarrow \frac{|\mathcal{D}_{\text{train}}^i|^{1/\tau}}{\sum_j |\mathcal{D}_{\text{train}}^j|^{1/\tau}}$ 
4: while  $\theta_t$  not converged (every  $k$  steps) do
5:    $\nabla_t \leftarrow 0, S \leftarrow \emptyset$  // gradient cache
6:   Sample val batch  $\{d_{\text{val}}^j\}_{j=1}^N \sim \mathcal{D}_{\text{val}}$ 
7:   for  $i = 1, \dots, M$  do
8:      $(x, y) \sim \mathcal{D}_{\text{train}}^i$ 
9:      $\theta_{t+1}^i \leftarrow \text{Step}(\theta_t, \text{Opt}_t; x, y)$  //  $\times l$  steps
10:     $\Delta \mathcal{M}_j^i \leftarrow \mathcal{M}(\theta_{t+1}^i; d_{\text{val}}^j) - \mathcal{M}(\theta_t; d_{\text{val}}^j)$ 
11:     $\mathcal{I}^i \leftarrow \frac{1}{N} \sum_j \Delta \mathcal{M}_j^i$ 
12:     $\nabla_t \leftarrow \nabla_t + \mathcal{I}^i (\theta_{t+1}^i - \theta_t), S \leftarrow S \cup \mathcal{I}^i$ 
13:   end for
14:    $\bar{\nabla}_t \leftarrow \nabla_t / S$ 
15:    $\theta_{t+1} \leftarrow \theta_t + \alpha \bar{\nabla}_t$  // reptile meta-update
16:    $\text{Opt}_{t+1} \leftarrow \text{StateUpdate}(\text{Opt}_t, \bar{\nabla}_t)$ 
17:    $d_\psi \leftarrow \sum_{i=1}^M P_D(i; \psi) \mathcal{I}^i \nabla_\psi \log P_D(i; \psi)$ 
18:    $\psi \leftarrow \text{GradientUpdate}(\psi, d_\psi)$ 
19: end while

```

3.3 Optimization for compute and scalability

To reuse gradients across datasets and reduce computation, we perform Reptile-style first-order meta-updates (Nichol et al., 2018). For each training dataset $\mathcal{D}_{\text{train}}^i$ we take l inner steps from the current initialization θ_t (step size η_t), producing θ_{t+1}^i , and compute an influence score $\mathcal{I}^i$. We convert scores to a sampling distribution via softmax, $p_i = \exp(\mathcal{I}^i / \tau) / \sum_j \exp(\mathcal{I}^j / \tau)$, and form the weighted Reptile update

$$\bar{\theta}_{t+1} = \sum_{i=1}^M p_i \theta_{t+1}^i \quad (5)$$

$$\theta_{t+1} = \theta_t + \alpha (\bar{\theta}_{t+1} - \theta_t)$$

with Reptile rate $\alpha = \eta_t$. Using this procedure we need only a single copy of parameter gradients and optimizer states (first/second moments), which substantially reduces memory overheads.

When M is large (e.g., many domains or languages) computing $\mathcal{I}^i$ for every i each iteration is costly. We therefore update the scorer ψ on a uniform random subsample $S \subset \{\mathcal{D}_{\text{train}}^i\}_{i=1}^M$ with $|S| = k < M$. Restricting the policy to S yields the conditional categorical

$$P_D(i | S; \psi) = \begin{cases} \frac{P_D(i; \psi)}{\sum_{j \in S} P_D(j; \psi)}, & i \in S, \\ 0, & i \notin S, \end{cases} \quad (6)$$

$$P_D(i; \psi) = \frac{\exp(\psi_i)}{\sum_{j=1}^M \exp(\psi_j)}.$$

We then compute the scorer gradient on S as

$$d_\psi = \sum_{i \in S} P_D(i; \psi) I^i \nabla_\psi \log P_D(i | S; \psi). \quad (7)$$

This estimator is unbiased for the *conditional* objective over S (by the policy-gradient identity) but biased with respect to the full-objective gradient over all M datasets. In practice, choosing $k < M$ yields large reductions in per-iteration compute and memory while still improving performance (see Fig. 8 and Section 4).

4 Experimental Setup

In this section, we detail the benchmarks and experimental setup used to test the domain adaptation of text retrievers under different sampling strategies.

BEIR: We start with a controlled setup where in-domain train, dev, and test retrieval datasets are available, though their sizes vary across domains. We train on all seven BEIR datasets that provide training splits: MS MARCO, NQ, FEVER, FiQA, HotpotQA, SciFact, and NFCorpus. Since only a subset of BEIR-15 datasets provide both official dev and test splits, evaluation and hyperparameter selection are performed on the five datasets where both are available: FEVER, Quora, FiQA, HotpotQA, and DBPedia, following the splits released by the BEIR authors. NFCorpus is excluded from the target datasets due to noise, as many level-1 relevant passages are in fact irrelevant. Our biencoder models are initialized with a pretrained *roberta-base* model (Liu et al., 2019), and trained for 2 epochs with InfoNCE loss (van den Oord et al., 2018) as influence metric $\mathcal{M}$ with cross-batch negatives. The scorer network is warmed up for 50 steps and updated every 50 training steps.

Multilingual Long Document Retrieval (MLDR): We next examine the realistic setting of Multilingual Long Document Retrieval, using the BGE-M3 corpus originally employed to train the *bge-m3-dense* model. Adaptation is performed on the MLDR-13 development set, with evaluation on its corresponding test set. The biencoder model is initialized from a 568M parameter *bge-m3-unsupervised* checkpoint. We treat each language as a domain and sample proportionally across all datasets in a language. The scorer network is warmed up for 500 steps and updated every 250 training steps. We use the *m3-kd-distill* loss from BGE-M3, with cross-batch negatives and 8 hard negatives, as our influence metric.

Sentence-Transformers Embedding Dataset:

Finally, we consider a more challenging scenario where train and test datasets span diverse domains, using the publicly available *sentence-transformers embedding dataset*, originally used to train the *all-MiniLM-L6-v2* model. The training corpus contains 440 million parallel sentences drawn from 32 datasets. Its data configuration includes carefully tuned sampling weights, referred to as *Expert* initialization, designed to optimize performance. To align with our target domain, we excluded the Reddit comments dataset due to its large size and the CodeSearchNet dataset, since our target evaluation tasks are not code-related. Our biencoder models are initialized with the pretrained *MiniLM-L6-H384-uncased* model, and we optimize performance jointly across all BEIR-5 dev sets. The scorer network is warmed up for 500 steps and updated every 250 steps. Following the toy setting, we use the standard InfoNCE loss as the influence metric with cross-batch negatives. Additional hyperparameters are listed in Appendix A.

Baselines: We compare our sampling algorithm against four categories of baselines: (i) static sampling methods (Temperature, Cooldown), (ii) a universal generalization approach (DoReMi), (iii) gradient-based task-adaptive sampling methods (MultiDDS, DoGE), and (iv) a cluster-level, task-adaptive importance-sampling method (CRISP). For all baselines, we follow the training guidelines recommended in their respective papers, with hyperparameter details provided in Appendix A. We evaluate statistical significance using a paired t-test comparing the two highest-performing models in each setting. For CRISP, we construct clusters in powers of 32^x : $x \in 1, 2$ for BEIR-train, $x \in 1, 2, 3$ for Sentence-Transformers, and $x \in 1, 2$ per language for BGE-M3.

5 Results And Analysis

In this section, we aim to answer the following research questions: (1) Does learning a dynamically evolving sampling distribution through influence measures lead to superior adaptation on the test set? (2) Does the influence-based scorer capture additional insights beyond domain similarity between training and dev sets? (3) In a diverse domain setting, how reliable are influence-based approaches compared to gradient-based methods?




Figure 3: FEVER training set sampling trajectories for different initialization temperatures using MultiDDS, Inf-DDS, and learned weights from baseline methods (optimized on FEVER dev set).




Figure 4: (a) Heatmap showing Z-score (row) normalized performance correlations between train and test splits across BEIR datasets. (b) Domain weights learned by Inf-DDS during optimization for each target domain.

5.1 Main Results

Domain adaptation on BEIR: Our initial findings demonstrates that mere domain similarity alone does not consistently lead to enhanced performance in text retrieval setting. Fig. 4(a) illustrates this by showing the normalized performance correlation between train/test sets without adaptive sampling. Notably, target datasets such as FEVER, HotpotQA, and FiQA benefit not only from their corresponding training domains but also from MS MARCO and $\tau = 1$ sampling, indicating that *effective retriever improvement requires going beyond domain similarity* highlighting the need for adaptive sampling strategies designed to optimize downstream performance.

To investigate this, we adapt a sampling distribution to the FEVER development set. As Inf-DDS relies on single-shot optimization, it is sensitive to the *initial distribution*, prompting us to assess multiple initializations to ensure robustness. As shown in Fig. 5(a), Inf-DDS consistently outperforms all baselines when initialized with $\tau = 0.3$, and remains competitive with CRISP using 32 clusters. Fig. 3 provides a comparison of sampling

trajectories, highlighting that MultiDDS exhibits unstable and inconsistent dynamics across varying temperatures, Inf-DDS produces stable behavior, consistently prioritizing FEVER and MS MARCO (Fig. 4(b)), in line with CRISP and DoReMi. This stability underscores the reliability of Inf-DDS in effectively aligning sampling strategies with measurable downstream performance improvements. Full results, including statistical significance tests, are presented in Table 7 in the Appendix.

To further validate our approach, we conduct joint optimization over the BEIR-5 dev sets for generalization. As shown in Fig. 5(b), Inf-DDS outperforms static sampling in all temperature initializations and surpasses MultiDDS in 2 out of 3 scenarios. While MultiDDS shows more improvements when initialized with $\tau = 1$, it underperforms Inf-DDS’ best score by 0.93 points. Table 9 in the Appendix presents the full results, including tests for statistical significance.

Multilingual Adaptation on MLDR: Multilingual retrieval introduces distinct challenges stemming from the substantial variability in language




Figure 5: Sampling probability initialization vs. Avg. NDCG@10 when training on the BEIR-7 train set under two optimization settings: (a) optimized on the FEVER dev set and evaluated on the FEVER test set, and (b) jointly optimized on the BEIR-5 dev sets and evaluated on the corresponding test sets. Adaptive sampling methods, Inf-DDS and MultiDDS, are compared under different initialization strategies.

resources and heterogeneous domain distributions. We assess how Inf-DDS and MultiDDS tackle these challenges by implicitly harmonizing data from high and low resource languages while leveraging cross-lingual relatedness without the need for explicit supervision. Specifically, we investigate two key questions: (1) Can Inf-DDS automatically up-sample underrepresented languages within a shared multilingual corpus to improve performance? (2) How critical is sampling high-resource languages when optimizing for multiple languages?




Figure 6: Avg. NDCG@10 scores on the MLDR-13 language test collection using BGE-M3 training data. Optimization of the scorer is done jointly on the 13 development sets.

We optimize the *bge-m3-unsupervised* model and scorer on the full MLDR-13 development set to maximize performance across all 13 languages. As shown in Fig. 6, starting from the same initial sampling weights as *bge-m3-dense*, Inf-DDS improves upon this baseline by +5.03 NDCG@10 points, while DoReMi achieves a +4.76-point gain. Inf-DDS also achieves the highest per-language performance in 7 of 13 languages. Full results, including statistical significance tests, are presented in Table 11 in the Appendix.

Sampling trajectories for each language are pre-

sented in Fig. 17 (Appendix). Interestingly, the sampling weights for English and Chinese drop substantially from their high initial values, yet performance remains comparable to *bge-m3-dense*, likely due to their dominant presence in *bge-m3-unsupervised* training (66.4% of total). This indicates that high-resource languages need little supervised data, demonstrating dynamic sampling’s ability to upweight low-resource languages while avoiding overfitting on dominant ones. Notably, Swahili’s sampling weight rises significantly despite minimal linguistic overlap with the development set; even with over half the training data sampled from Swahili, the model outperforms the baseline, though the cause of this effect remains unclear.

Sentence-Transformers Embedding Dataset:

This diverse training corpus includes over 440 Million query-positive passage pairs spanning 32 domains. We use the same BEIR-5 development sets from the toy setting for optimization, as they have minimal overlap with the training data. This experiment addresses two key questions: (1) Does dynamic sampling remain effective with high domain diversity? (2) Can our algorithm further improve performance when a strong initial sampling distribution is available? The extensive domain coverage significantly increases the complexity of the adaptation problem.

Starting from a uniform initialization, Inf-DDS achieves performance only 0.22 points below the off-the-shelf Sentence-Transformers Expert model (*all-MiniLM-L6-v2*), nearly optimal—and yields a 1.83 point gain over the uniform baseline. In con-




Figure 7: Average NDCG@10 on the BEIR-5 test collection using Sentence-Transformers training data with uniform (a) and *expert* initialization (b). The scorer is optimized jointly on the development sets.

trast, gradient-based methods such as MultiDDS and DoGE fail to make any gains. When we re-run the experiment starting from Expert weights, Inf-DDS still produces an additional 0.94 point improvement, demonstrating its ability to refine even strong initial distributions. DoReMi attains a larger 1.25 point gain but requires $3\times$ the compute (Fig. 8). Fig. 15 shows the evolving sampling distributions under Inf-DDS, illustrating that Expert weights can be further improved by dynamically adapting dataset sampling.

5.2 Discussion

Computational Overheads: Inf-DDS computes exact influence scores for each training dataset using online proxy models. While this introduces additional overhead, retrieval models are typically small, making the tradeoff worthwhile given the performance gains. Fig. 8 compares the training time of *all-MiniLM-L6-v2* across different sampling strategies. We also provide a computational breakdown of the time spent on model optimization (θ) and sampler optimization (ψ) in Table 1. Although Inf-DDS is slightly slower than CRISP, MultiDDS, and static sampling, it consistently achieves superior performance. To reduce memory usage, Inf-DDS stores only a single set of intermediate gradients and optimizer states during influence computation, which are efficiently reused in the weighted Reptile update.




Figure 8: Comparison of approximate GPU Hours for training on Sent-Trans embedding data.

Setting	Model (θ)	Sampler (ψ)
Sent-Trans	1258	178
MLDR	916	408

Table 1: GPU-hour breakdown of Inf-DDS for θ training and ψ influence updates.

Effect of initialization: Choosing an effective initialization for sampling is challenging but can substantially impact Inf-DDS’s performance (Figs. 5(a), 5(b) and 7). While the algorithm does not always reach globally optimal sampling weights, starting from a reasonable initialization and updating the weights consistently yields gains. We do not explore heuristics for selecting initial weights, but experiments with standard static initializations show that, although no single choice is universally best, reasonably good initializations generally perform well.

Effect of Reptile Updates and Update Steps: We conduct an ablation study on BEIR to evaluate the contribution of the meta-learning component, with results reported in Table 12 (Appendix). Disabling Reptile updates leads to only marginal performance changes, suggesting that the primary gains stem from dynamic sampling rather than meta-learning itself. Nevertheless, Reptile remains beneficial in practice, as it reduces computational overhead by enabling reuse of intermediate computations. We further analyze the effect of the number of update steps (l) required to obtain reliable influence estimates on the FEVER and FiQA datasets under the BEIR setting. As shown in Table 13, using 3–5 update steps provides a good trade-off, yielding strong performance while avoiding unnecessary computation.

Dev/Test Overlap Analysis: Adaptive data selection methods such as DoGE, DDS, CRISP, and Inf-DDS use development sets to compute rewards or influence estimates, raising potential concerns about dev–test leakage. While our main experiments follow the standard train, dev, and test splits from the original benchmarks, we additionally verify that this setup does not introduce test bias.

We first train models using the original MLDR and BEIR training splits. We then perform repeated train–dev resampling by merging the training and development sets and sampling five new training sets matching the original training size, while always evaluating on the original test sets. Results in Tables 14 and 15 show that models trained on mixed train–dev folds perform comparably to mod-

els trained only on the original training data (e.g., MLDR mixed average 39.6 ± 0.3 vs. train baseline 40.0; BEIR mixed average 46.3 ± 0.2 vs. train baseline 45.9), indicating that the development sets do not leak information into the test sets.

Relation between Influence and Gradients:

Stochastic gradient updates optimize the model parameters with respect to the loss $\mathcal{L}$, while influence measures their effect on the target metric $\mathcal{M}$. When $\mathcal{L}$ and $\mathcal{M}$ are closely aligned, a first-order approximation shows that the single-step change in $\mathcal{M}$ depends on the alignment between the development-set gradient ∇_{dev} and the training gradient ∇_{train} . This corresponds to the gradient-alignment reward used by methods such as DDS/MultiDDS, making it a first-order approximation to our single-step influence reward. Inf-DDS instead computes influence from the realized change in $\mathcal{M}$ after multiple ($l > 1$) inner optimization steps, capturing the cumulative effect of the full proxy trajectory rather than a local first-order approximation. In high-dimensional optimization landscapes, gradient alignment at a single point may not accurately reflect the eventual effect of an update, particularly as the optimization trajectory moves away from that point. Influence-based rewards instead evaluate training data based on the resulting change in the target metric after the updates. We additionally observe empirically (Section 5 and Fig. 3) that gradient-based approaches exhibit higher variance and less stable sampling trajectories in practice.

6 Conclusion

In this work, we present a comprehensive analysis of adapting text retrievers to target domains using influence-guided dynamic data sampling (Inf-DDS). Our approach parametrizes the sampling distribution with scorer parameters ψ and performs bi-level optimization, jointly updating model parameters θ and scorer parameters ψ using influence scores as rewards. Across multiple benchmarks and baselines spanning diverse domains, Inf-DDS produces stable sampling trajectories and consistently matches or outperforms proxy-model and gradient-based approaches while remaining computationally efficient. Although the algorithm does not always converge to the global optimum, it reliably delivers substantial improvements from reasonable initializations. We further analyze why gradient-based signals can mislead optimization, demonstrating that influence-based rewards provide a more robust

estimate of the best attainable minima. For future work, we aim to investigate improved initialization strategies and optimization techniques for the parameterized scorer distribution ψ .

Limitations

Our experiments primarily focus on adapting text retrievers using a fixed set of training and development datasets. An important future direction is studying whether dynamic sampling can further adapt already fine-tuned models to new domains. Although our method consistently improves performance, it does not always converge to the optimal sampling distribution. We observe that strong initializations can help, but better optimization strategies are still needed to reach high-performing distributions efficiently. Another open question is identifying when dynamic sampling is most effective, particularly in settings where domains are highly conflated and difficult to reweight.

References

- Juhan Bae, Nathan Huyen Ng, Alston Lo, Marzyeh Ghassemi, and Roger Baker Grosse. 2022. [If influence functions are the answer, then what is the question?](#) In *Advances in Neural Information Processing Systems*.
- Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. 2024. [Llm2vec: Large language models are secretly powerful text encoders](#). *CoRR*, abs/2404.05961.
- Luiz Henrique Bonifacio, Israel Campiotti, Roberto A. Lotufo, and Rodrigo Nogueira. 2021. [mmarco: A multilingual version of MS MARCO passage ranking dataset](#). *CoRR*, abs/2108.13897.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Yihan Cao, Yanbin Kang, Chi Wang, and Lichao Sun. 2024. [Instruction mining: Instruction data selection for tuning large language models](#). In *First Conference on Language Modeling*.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [BGE m3-embedding](#):

- Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *CoRR*, abs/2402.03216.
- Dami Choi, Derrick Xin, Hamid Dadkhahi, Justin Gilmer, Ankush Garg, Orhan Firat, Chih-Kuan Yeh, Andrew M. Dai, and Behrooz Ghorbani. 2023. [Order matters in the presence of dataset imbalance for multilingual learning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Hyung Won Chung, Xavier Garcia, Adam Roberts, Yi Tay, Orhan Firat, Sharan Narang, and Noah Constant. 2023. [Unimax: Fairer and more effective language sampling for large-scale multilingual pre-training](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Logan Engstrom, Axel Feldmann, and Aleksander Madry. 2024. Dsdm: model-aware dataset selection with datamodels. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Simin Fan, Matteo Pagliardini, and Martin Jaggi. 2024. Doge: domain reweighting with generalization estimation. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- David Grangier, Pierre Ablyn, and Awni Hannun. 2025a. [Adaptive training distributions with scalable online bilevel optimization](#). In *Transactions on Machine Learning Research (TMLR)*.
- David Grangier, Simin Fan, Skyler Seto, and Pierre Ablyn. 2025b. [Task-adaptive pretrained language models via clustered-importance sampling](#). In *The Thirteenth International Conference on Learning Representations*.
- Roger B. Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, Evan Hubinger, Kamile Lukosiute, Karina Nguyen, Nicholas Joseph, Sam McCandlish, Jared Kaplan, and Samuel R. Bowman. 2023. [Studying large language model generalization with influence functions](#). *CoRR*, abs/2308.03296.
- Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu, Qiaoqiao She, Xuan Liu, Tian Wu, and Haifeng Wang. 2018. [DuReader: a Chinese machine reading comprehension dataset from real-world applications](#). In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 37–46, Melbourne, Australia. Association for Computational Linguistics.
- Andrew Ilyas, Sung Min Park, Logan Engstrom, Guillaume Leclerc, and Aleksander Madry. 2022. [Data-models: Predicting predictions from training data](#). *CoRR*, abs/2202.00622.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Trans. Mach. Learn. Res.*, 2022.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Mi-Young Kim, Juliano Rabelo, Randy Goebel, Masaharu Yoshioka, Yoshinobu Kano, and Ken Satoh. 2022. [Collee 2022 summary: Methods for legal document retrieval and entailment](#). In *New Frontiers in Artificial Intelligence: JSAI-IsAI 2022 Workshop, JURISIN 2022, and JSAI 2022 International Session, Kyoto, Japan, June 12–17, 2022, Revised Selected Papers*, page 51–67, Berlin, Heidelberg. Springer-Verlag.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1885–1894. PMLR.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.

- Haitao Li, Yunqiu Shao, Yueyue Wu, Qingyao Ai, Yixiao Ma, and Yiqun Liu. 2024a. [Lecardv2: A large-scale chinese legal case retrieval dataset](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, page 2251–2260, New York, NY, USA. Association for Computing Machinery.
- Tianjian Li, Haoran Xu, Weiting Tan, Kenton Murray, and Daniel Khashabi. 2024b. [Upsample or up-weight? balanced training on heavily imbalanced datasets](#). *CoRR*, abs/2410.04579.
- Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. 2025. [Regmix: Data mixture as regression for language model pre-training](#). In *The Thirteenth International Conference on Learning Representations*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Robert C. Moore and William Lewis. 2010. [Intelligent selection of language model training data](#). In *Proceedings of the ACL 2010 Conference Short Papers*, pages 220–224, Uppsala, Sweden. Association for Computational Linguistics.
- 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 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.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. [On first-order meta-learning algorithms](#). *CoRR*, abs/1803.02999.
- Mahdi Nikdan, Vincent Cohen-Addad, Dan Alistarh, and Vahab Mirrokni. 2025. Efficient data selection at scale via influence distillation. *arXiv preprint arXiv:2505.19051*.
- Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. 2023. [Trak: attributing model behavior at scale](#). In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org.
- Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. [Estimating training data influence by tracing gradient descent](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogenous benchmark for zero-shot evaluation of information retrieval models](#). *CoRR*, abs/2104.08663.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation learning with contrastive predictive coding](#). *CoRR*, abs/1807.03748.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2023. [SimLM: Pre-training with representation bottleneck for dense passage retrieval](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2244–2258, Toronto, Canada. Association for Computational Linguistics.
- Xinyi Wang, Hieu Pham, Paul Michel, Antonios Anastopoulos, Jaime G. Carbonell, and Graham Neubig. 2020a. [Optimizing data usage via differentiable rewards](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9983–9995. PMLR.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020b. [Balancing training for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8526–8537, Online. Association for Computational Linguistics.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Mach. Learn.*, 8:229–256.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. 2024. [LESS: selecting influential data for targeted instruction tuning](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V Le, Tengyu Ma, and Adams Wei Yu. 2023a. [Doremi: Optimizing data mixtures speeds up language model](#)

- [pretraining](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy Liang. 2023b. [Data selection for language models via importance resampling](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Xiaohui Xie, Qian Dong, Bingning Wang, Feiyang Lv, Ting Yao, Weinan Gan, Zhijing Wu, Xiangsheng Li, Haitao Li, Yiqun Liu, and Jin Ma. 2023c. [T2ranking: A large-scale chinese benchmark for passage ranking](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 2681–2690, New York, NY, USA. Association for Computing Machinery.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Zichun Yu, Spandan Das, and Chenyan Xiong. 2024. [MATES: Model-aware data selection for efficient pretraining with data influence models](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Chi Zhang, Huaping Zhong, Kuan Zhang, Chengliang Chai, Rui Wang, Xinlin Zhuang, Tianyi Bai, Qiu Jiantao, Lei Cao, Ju Fan, Ye Yuan, Guoren Wang, and Conghui He. 2025. [Harnessing diversity for important data selection in pretraining large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Sheng Zhang, Xin Zhang, Hui Wang, Lixiang Guo, and Shanshan Liu. 2018. [Multi-scale attentive interaction networks for chinese medical question answer selection](#). *IEEE Access*, 6:74061–74071.
- Xinyu Zhang, Xueguang Ma, Peng Shi, and Jimmy Lin. 2021. [Mr. TyDi: A multi-lingual benchmark for dense retrieval](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 127–137, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023. [MIRACL: A multilingual retrieval dataset covering 18 diverse languages](#). *Transactions of the Association for Computational Linguistics*, 11:1114–1131.
- Haotian Zhou, Tingkai Liu, Qianli Ma, Yufeng Zhang, Jianbo Yuan, Pengfei Liu, Yang You, and Hongxia Yang. 2025. [DavIR: Data selection via implicit reward for large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9220–9237, Vienna, Austria. Association for Computational Linguistics.
- Zhi Zhou, Jiang-Xin Shi, Peng-Xiao Song, Xiaowen Yang, Yi-Xuan Jin, Lan-Zhe Guo, and Yufeng Li. 2024. [Lawgpt: A chinese legal knowledge-enhanced large language model](#). *CoRR*, abs/2406.04614.

A Hyperparameters

A.1 BEIR

For our experiments, we initialize our bi-encoder with *roberta-base* model, we set the initial learning rate to $2e-5$, and train with mixed-precision (BF16) enabled. We used an in-batch negative sampling strategy with a temperature of 0.02, normalizing representations before contrastive scoring. Both training and evaluation batch sizes were 256 examples on a single NVIDIA-A100 80GB GPU. For scorer updates, we also use a batch size of 256. We use a linear learning rate decay with a warmup of 250 steps, and trained for a total of 7,000 steps. We use the standard InfoNCE (van den Oord et al., 2018) as our loss and our influence metric $\mathcal{M}$. Queries and passages were truncated to maximum lengths of 64 and 256 tokens, respectively, with CLS-token pooling for sentence embeddings. We further update our scorer every 50 steps after a 50-step scorer warmup (we don’t train the scorer during the warmup period). For reptile updates, we use α based on the current learning rate η_t of the learning rate schedule. We do checkpoint selection by selecting the best checkpoint on the dev set and report the corresponding numbers on the final test set. For DoReMi, we train our reference models using $\tau = 1$ sampling and proxy models starting with $\tau = \infty$ initialization. For DoGE, too, we train our proxy models starting with $\tau = \infty$ initialization. For CRISP, we limit the clusters to 32^x : $x \in 1, 2, 3$ since the size of the total dataset is only 800k instances and use *all-MiniLM-L6-v2* embedding model for clustering.

A.2 Sentence Transformers Embedding

We initialize our bi-encoder with *nreimers/MiniLM-L6-H384-uncased* model, we set the initial learning rate to $2e-5$, and train with mixed-precision (BF16) enabled. We used an in-batch and cross-device negative sampling strategy with a temperature of 0.02, normalizing representations before contrastive scoring. Training batch sizes are set to 2000, and evaluation batch sizes are 256 examples per GPU and use 8 NVIDIA-A100 80GB GPUs. We use the standard InfoNCE (van den Oord et al., 2018) as our loss and our influence metric $\mathcal{M}$. We also use 1-hard negative when using MSMarco in the training set. For scorer updates, we use a batch size of 256 per dataset. We use a linear learning rate decay with a warmup of 1000 steps, and trained for a total of 150k steps. Queries and passages

were truncated to maximum lengths of 64 and 256 tokens, respectively, with CLS-token pooling for sentence embeddings. We further update our scorer every 250 steps after a 500-step scorer warmup (we don’t train the scorer during the warmup period). For reptile updates, we use α based on the current learning rate η_t of the learning rate schedule. For DoReMi, we train our reference models using $\tau = \infty$ sampling when comparing for uniform initialization $\tau = \text{expert}$ for expert initialization and proxy models starting with $\tau = \infty$ initialization. For DoGE, too, we train our proxy models starting with $\tau = \infty$ initialization sampling when comparing for uniform initialization $\tau = \text{expert}$ for expert initialization. For CRISP, we limit the clusters to 32^x : $x \in 1, 2, 3$ since the size of the total dataset is only 440M instances, and use *all-MiniLM-L6-v2* embedding model for clustering.

A.3 Multilingual Long Document Retrieval

We initialize our bi-encoder with *BAAI/bge-m3-unsupervised* model, we set the initial learning rate to $2e-5$, and train with mixed-precision (BF16) enabled. We used an in-batch negative, cross-device negatives, and 8 hard negatives during training with a temperature of 0.02, normalizing representations before contrastive scoring. We use the *bge-m3-kd-distil* loss by BGE-M3 as our training loss and InfoNCE with hard negatives as our influence metric $\mathcal{M}$. Training batch sizes are set to 5, and evaluation batch sizes are 5 examples per GPU and use 8 NVIDIA-A100 80GB GPUs. For scorer updates, we use a batch size of 4 per dataset. We use a linear learning rate decay with a warmup of 1000 steps, and trained for a total of 10k steps. Queries and passages were truncated to maximum lengths of 512 and 8192 tokens, respectively, with CLS-token pooling for sentence embeddings. We further update our scorer every 250 steps after a 500-step scorer warmup (we don’t train the scorer during the warmup period). We employ subsampling as detailed in Section 3.3 with $k = 8$. We experiment with both enabling and disabling Reptile updates, and observe better performance when Reptile updates are turned off in the MLDR-13 setting. We use reptile update α based on the current learning rate η_t of the learning rate schedule. For DoReMi, we train our reference models using $\tau = \infty$ sampling and proxy models starting with $\tau = \infty$ initialization. For DoGE, too, we train our proxy models starting with $\tau = \infty$ initialization. For CRISP, we limit the clusters to 32^x : $x \in 1, 2$ per

language since the size of the total dataset is only 1.5M instances and use *paraphrase-multilingual-MiniLM-L12-v2* embedding model for clustering.

B Datasets

B.1 BEIR

Table 2: BEIR Train Datasets

Dataset	#Qrels
MS MARCO	499,184
NFCorpus	2,590
NQ	100,231
HotpotQA	85,000
FiQA	5,500
Fever	109,810
SciFact	809
Total	803,124

Table 3: BEIR Dev Datasets

Dataset	#Qrels
HotpotQA	10,894
FiQA	1,238
Quora	7,626
DBpedia	5,673
Fever	8,079
Total	40,947

B.2 Sentence Transformers Embedding

Dataset

The all-MiniLM-L6-v2 model was fine-tuned with a self-supervised contrastive objective over a concatenation of diverse, publicly available sentence-pair corpora. These include user-generated content such as paired Reddit comments and Q&A threads from Stack Exchange and Yahoo Answers; scientific citation pairs drawn from the S2ORC and SPECTER datasets; question answering benchmarks like PAQ, MS MARCO, Natural-Questions, SearchQA, SQuAD 2.0, and TriviaQA; paraphrase and duplicate-question collections from WikiAnswers, Quora Question Triplets, and the AllNLI (SNLI+MultiNLI) corpus; multimodal captions from COCO and Flickr30k; code-search examples; and specialized text-compression and instructional corpora such as Simple Wikipedia, Wikihow, Altlex, and explicit sentence-compression datasets. Full dataset statistics are provided in Table 4.

B.3 BGE-m3 Multilingual

For English, bge-m3 fine-tuning dataset includes 8 datasets, including HotpotQA (Yang et al., 2018), TriviaQA (Joshi et al., 2017), NQ (Kwiatkowski et al., 2019), MS MARCO (Nguyen et al., 2016), COLIEE (Kim et al., 2022), PubMedQA (Jin et al., 2019), SQuAD (Rajpurkar et al., 2016), and SimCSE (Gao et al., 2021). For Chinese, it includes 7 datasets, DuReader (He et al., 2018), mMARCO-ZH (Bonifacio et al., 2021), T2-Ranking (Xie et al., 2023c), LawGPT (Zhou et al., 2024), CMedQAv2 (Zhang et al., 2018), NLIzh2, and LeCaRDv2 (Li et al., 2024a). It also includes training data for other languages from Mr. Tydi (Zhang et al., 2021), MIR-ACL (Zhang et al., 2023) and train sets of MLDR. Full dataset statistics are provided in Table 6.

C Additional results

C.1 BEIR

To further validate our approach, we conduct individual optimizations on each of the BEIR-5 development sets and report the results in Table 8 and Figure 9. As shown, Inf-DDS consistently matches or outperforms static sampling across all datasets. While it does not achieve the highest performance in every domain compared to MultiDDS, Inf-DDS surpasses proportional sampling by an average margin of 2.24 points and even outperforms MultiDDS on average. The corresponding sampling trajectories are provided in the Appendix D.

C.2 Sentence Transformers Embedding

Dataset

Here in Table 10 we report Average NDCG@10 numbers on the BEIR-5 test collection when training *all-MiniLM-L6-v2* using Sentence-Transformers data, when optimization of scorer is done jointly on the BEIR-5 development sets. The corresponding sampling trajectories are shown in Figs. 15 and 16.

C.3 MLDR

Here in Table 11 we report Average NDCG@10 numbers on MLDR-13 test set when training with BGE-M3 data, when optimization of scorer is done jointly on the MLDR-13 development sets.

C.4 Ablation

Effect of reptile updates: We perform an ablation study to determine whether the observed performance gains arise solely from dynamically

Dataset	Count	Proportional Sampling %	Weight	Expert Sampling %
S2ORC Citation pairs (Abstracts)	116,288,806	26.42%	123	3.07%
WikiAnswers Duplicate question pairs	77,427,422	17.60%	123	3.07%
Amazon QA Pairs	2,448,839	0.56%	247	6.16%
PAQ (Question, Answer) pairs	64,371,441	14.63%	123	3.07%
S2ORC Citation pairs (Titles)	52,603,982	11.96%	123	3.07%
S2ORC (Title, Abstract)	41,769,185	9.49%	123	3.07%
Stack Exchange (Title, Body) pairs	25,316,456	5.75%	565	14.09%
Stack Exchange (Title+Body, Answer) pairs	21,396,559	4.86%	17	0.42%
Stack Exchange (Title, Answer) pairs	21,396,559	4.86%	373	9.31%
Stack Exchange Math	2,218,989	0.50%	166	4.14%
MS MARCO triplets	9,144,553	2.08%	247	6.16%
GOOQAQ: Open QA with Diverse Answer Types	3,012,496	0.68%	247	6.16%
Yahoo Answers (Title, Answer)	1,198,260	0.27%	247	6.16%
COCO Image captions	828,395	0.19%	1	0.02%
SPECTER citation triplets	684,100	0.16%	84	2.10%
Yahoo Answers (Question, Answer)	681,164	0.15%	169	4.21%
Yahoo Answers (Title, Question)	659,896	0.15%	163	4.07%
SearchQA	582,261	0.13%	144	3.59%
Eli5	325,475	0.07%	81	2.02%
Flickr 30k	317,695	0.07%	1	0.02%
Stack Exchange Duplicate questions (titles)	304,525	0.07%	26	0.65%
AllNLI (SNLI and MultiNLI)	277,230	0.06%	69	1.72%
Stack Exchange Duplicate questions (bodies)	250,519	0.06%	21	0.52%
Stack Exchange Duplicate questions (titles+bodies)	250,460	0.06%	21	0.52%
Sentence Compression	180,000	0.04%	45	1.12%
Wikihow	128,542	0.03%	32	0.80%
Altlex	112,696	0.03%	28	0.70%
Quora Question Triplets	103,663	0.02%	26	0.65%
Simple Wikipedia	102,225	0.02%	26	0.65%
Natural Questions (NQ)	100,231	0.02%	25	0.62%
SQuAD2.0	87,599	0.02%	22	0.55%
TriviaQA	73,346	0.02%	19	0.47%
Total	440,096,511	100.00%	4009	100.00%

Table 4: Training data provided by sentence transformers *all-MiniLM-L6-v2* showing dataset sizes, hand-picked weights, and normalized sampling percentages.

updating the sampling distribution or whether the meta-learning component of our updates also contributes. Table 12 compares downstream performance with Reptile updates enabled versus disabled. We observe only minor differences in performance, which suggests that most of the improvement can be attributed to dynamic sampling. However, because the Reptile meta-update reuses existing computations, it remains valuable for reducing overall computational overhead.

D Sampling trajectories

Language	Sampling (%)
Swahili	0.588
Farsi	0.588
Finnish	0.294
Indonesian	0.294
French	1.176
German	1.176
Korean	1.176
Spanish	1.176
Italian	1.176
Portuguese	1.176
Japanese	1.176
Bengali	0.294
Telugu	0.294
Thai	1.176
Russian	2.353
Hindi	1.176
Arabic	2.353
Chinese	23.529
English	58.824

Table 5: Language wise initialization probabilities for *bge-m3-dense* training.

Dataset	Lines	Sampling %	Dataset	Lines	Sampling %
MS MARCO	485,905	30.95%	Mr.TyDi/finnish	6,561	0.42%
MIRACL/fr	1,143	0.07%	Mr.TyDi/bengali	1,713	0.11%
MIRACL/zh	1,312	0.08%	Mr.TyDi/russian	5,366	0.34%
MIRACL/es	2,162	0.14%	Mr.TyDi/swahili	2,072	0.13%
MIRACL/ja	3,477	0.22%	Mr.TyDi/indonesian	4,902	0.31%
MIRACL/te	3,452	0.22%	Mr.TyDi/arabic	12,377	0.79%
MIRACL/en	2,863	0.18%	Mr.TyDi/english	3,547	0.23%
MIRACL/id	4,071	0.26%	Mr.TyDi/korean	1,295	0.08%
MIRACL/fa	2,107	0.13%	Mr.TyDi/japanese	3,697	0.24%
MIRACL/ko	868	0.06%	Mr.TyDi/telugu	3,880	0.25%
MIRACL/fi	2,897	0.18%	Mr.TyDi/thai	3,319	0.21%
MIRACL/th	2,972	0.19%	T2Ranking	90,467	5.76%
MIRACL/bn	1,631	0.10%	en_NLI/nli_for_simcse	274,951	17.51%
MIRACL/ru	4,683	0.30%	Law-Medical/colliee	463	0.03%
MIRACL/hi	1,169	0.07%	Law-Medical/law_gpt	500	0.03%
MIRACL/ar	3,495	0.22%	Law-Medical/lecardv2	591	0.04%
MIRACL/sw	1,901	0.12%	Law-Medical/pubmed_qa	500	0.03%
HotpotQA	84,516	5.38%	MLDR/hi	1,618	0.10%
mMARCO-zh/chinese	100,000	6.37%	MLDR/es	2,254	0.14%
NQ	58,568	3.73%	MLDR/ru	1,864	0.12%
zh_NLI/LCQMC	10,000	0.64%	MLDR/de	1,847	0.12%
zh_NLI/BQ	12,599	0.80%	MLDR/ja	2,262	0.14%
zh_NLI/STS-B	249	0.02%	MLDR/fr	1,608	0.10%
zh_NLI/afqmc	10,534	0.67%	MLDR/ar	1,817	0.12%
zh_NLI/ATEC	11,325	0.72%	MLDR/ko	2,198	0.14%
zh_NLI/QBQTC_v2	10,000	0.64%	MLDR/en	10,000	0.64%
zh_NLI/PAWSX	10,000	0.64%	MLDR/zh	10,000	0.64%
DuReader	80,416	5.12%	MLDR/pt	1,845	0.12%
cMedQAv2	50,000	3.18%	MLDR/it	2,151	0.14%
TriviaQA	60,315	3.84%	MLDR/th	1,970	0.13%
			SQuAD	87,599	5.58%
			Total	1,569,864	100.00%

Table 6: Training data provided by *bge-m3* showing dataset sizes and normalized sampling percentages.




Figure 9: Average NDCG@10 on the BEIR-5 test collection using BEIR-7 training data with $\tau = 1$. The scorer is optimized individually on the development sets.

Init.	Static Sampling	MultiDDS	Inf-DDS	Others
$\tau = 0.3^\dagger$	57.3	54.5	64.9*	-
$\tau = 1$	48.8	53.9	54.7	-
$\tau = 5$	46.5	53.2	55.2	-
$\tau = \infty$	50.1	55.8	57.5	-
Cooldown	-	-	-	50.5
DoReMi	-	-	-	48.6
DoGE	-	-	-	54.7
CRISP-32	-	-	-	59.0
CRISP-1024 [‡]	-	-	-	62.9

Table 7: NDCG@10 comparison of sampling strategies with varying initialization temperatures during optimization on the FEVER dev set. [†] and [‡] mark the two models compared in the paired significance test, and * indicates a statistically significant difference ($p < 0.05$). Additional results from other baselines are shown in the last column.

Sampling	Init.	Train Dataset	Dev Dataset	Average Test	DBpedia	FEVER	FiQA	HotpotQA	Quora
Static	$\tau = 1$	BEIR-7	-	48.7	29.5	48.8	29.0	52.4	84.0
MultiDDS	$\tau = 1$	BEIR-7	BEIR-1	50.4	27.4	53.9	33.3	53.8	83.7
Inf-DDS	$\tau = 1$	BEIR-7	BEIR-1	51.0	29.4	54.7	33.7	52.9	84.1
DoReMi	Proxy	BEIR-7	-	49.0	27.5	48.6	32.0	53.7	83.1
DoGE	Proxy	BEIR-7	BEIR-1	50.0	25.1	54.7	32.4	54.8	82.8
CRISP-32	Cluster IS	BEIR-7	BEIR-1	50.8	27.0	59.0	31.3	52.9	83.6
CRISP-1024	Cluster IS	BEIR-7	BEIR-1	51.1	24.8	62.9	31.8	52.6	83.3
Static	$\tau = \infty$	BEIR-7	-	48.9	27.2	50.1	31.7	52.8	82.8
MultiDDS	$\tau = \infty$	BEIR-7	BEIR-1	49.8	25.6	55.8	31.1	53.1	83.1
Inf-DDS	$\tau = \infty$	BEIR-7	BEIR-1	50.3	26.7	57.5	32.0	52.6	82.7

Table 8: Average NDCG@10 on BEIR-5 test collection, optimization of scorer is done individually on the development sets.

Sampling	Init	Train Dataset	Dev Dataset	Avg. Test	Dbpedia	Fever	Fiqu	Hotpotqa	Quora
Static	$\tau = 0.3$	BEIR-7	-	49.2	28.6	58.2	26.5	48.7	84.0
MultiDDS			BEIR-5	48.2	25.2	58.6	23.3	54.0	80.1
Inf-DDS [†]				50.3	27.7*	69.2*	23.6	48.6	82.3*
Static	$\tau = 1$		-	48.7	29.5	48.8	29.0	52.4	84.0
MultiDDS			BEIR-5	49.3	25.9	59.0	29.0	50.5	82.0
Inf-DDS				48.9	29.5	50.0	29.4	51.6	84.1
Static	$\tau = 5$		-	48.9	27.6	48.4	32.4	53.2	83.0
MultiDDS			BEIR-5	47.4	24.4	51.0	30.1	49.7	81.5
Inf-DDS				49.7	28.8	55.0	31.2	51.4	82.2
Static	$\tau = \infty$		-	48.9	27.2	50.1	31.7	52.8	82.8
MultiDDS			BEIR-5	47.3	22.7	52.4	30.4	50.1	81.1
Inf-DDS				50.2	26.8	59.6	32.1	50.4	82.2
DoReMi	Proxy	BEIR-7	-	49.0	27.5	48.6	32.0	53.7	83.1
DoGE [‡]	Proxy	BEIR-7	BEIR-5	50.4	26.7	61.3	30.4*	51.4*	82.2
CRISP-32	Cluster IS	BEIR-7	BEIR-5	50.0	26.2	57.2	31.0	52.3	83.5
CRISP-1024	Cluster IS	BEIR-7	BEIR-5	49.4	23.6	56.1	31.2	52.5	83.3

Table 9: Average NDCG@10 on the BEIR-5 test collections. Scorer optimization is performed jointly on the development sets. [†] and [‡] denote the two models compared in the paired significance test, and * indicates a statistically significant difference ($p < 0.05$). Additional baselines are listed below.

Sampling	Init.	Train Dataset	Dev Dataset	Average Test	DBpedia	FEVER	FiQA	HotpotQA	Quora
Static	$\tau = \infty$	Sent-Trans	-	49.0	31.7	49.0	31.8	44.6	49.0
MultiDDS	$\tau = \infty$	Sent-Trans	BEIR-5	40.2	22.3	43.9	24.8	26.5	83.9
InfDDS	$\tau = \infty$	Sent-Trans	BEIR-5	50.8	32.0	50.6	34.4	46.6	88.0
Cooldown	$5 \rightarrow 1$	Sent-Trans	-	47.6	28.6	49.9	33.7	38.1	87.6
DoReMi	Proxy	Sent-Trans	-	51.4	31.7	52.3	36.2	49.1	87.6
DoGE	Proxy	Sent-Trans	BEIR-5	49.5	29.9	57.4	31.2	42.6	86.3
CRISP (n=32)	Cluster IS	Sent-Trans	BEIR-5	47.4	28.3	58.1	27.9	38.0	84.9
CRISP (n=32 ²)	Cluster IS	Sent-Trans	BEIR-5	45.6	26.8	54.8	26.4	35.5	84.7
CRISP (n=32 ³)	Cluster IS	Sent-Trans	BEIR-5	48.5	28.4	61.9	28.5	39.5	84.2
Static (all-MiniLM-L6-v2) [‡]	Expert	Sent-Trans	-	51.0	32.3	51.9	36.9*	46.5*	87.6
MultiDDS	Expert	Sent-Trans	BEIR-5	42.5	22.8	37.2	29.2	28.3	86.3
InfDDS [†]	Expert	Sent-Trans	BEIR-5	52.0	32.4*	58.3*	36.1	45.7	87.4
DoReMi	Expert	Sent-Trans	-	52.3	32.5	56.3	36.2	48.7	87.7
DoGE	Expert	Sent-Trans	BEIR-5	50.5	30.0	59.6	33.3	43.1	86.4

Table 10: Average NDCG@10 on BEIR-5 test collection when training *all-MiniLM-L6-v2* using Sentence-Transformers data, optimization of scorer is done jointly on BEIR-5 development sets. [†] and [‡] mark the two models compared in the paired significance test, and * indicates a statistically significant difference ($p < 0.05$).

Model/ Sampling	Dev Dataset	Avg. Test	ar	de	en	es	fr	hi	it	ja	ko	pt	ru	th	zh
bge-m3 unsup.	-	37.0	30.8	38.4	34.2	61.4	53.1	23.9	43.7	33.3	24.9	59.7	45.8	18.6	13.7
bge-m3-dense	-	52.5	47.6	46.1	48.9	74.8	73.8	40.7	62.7	50.9	42.9	74.4	59.5	33.6	26.0
bge-m3-dense [◊]	-	52.5	48.4	46.7	46.7	76.3	74.3	40.0	61.7	49.1	41.0	74.3	60.8	36.0	26.7
MultiDDS	MLDR-13	56.7	56.0	53.9	50.0	79.1	76.3	44.1	66.0	56.6	49.3	78.6	62.3	37.3	27.3
Inf-DDS [†]	MLDR-13	57.5	58.3*	54.4	48.5	80.6*	77.9*	43.9	66.1	56.7*	49.9*	79.5	65.7*	39.0	26.6
Cooldown	-	52.7	51.5	50.2	48.5	75.6	72.6	33.1	61.7	48.4	41.3	75.5	61.4	37.8	27.7
DoReMi [‡]	-	57.2	54.8	54.4	50.7*	79.1	77.1	44.5*	66.6*	54.1	48.0	82.3*	64.5	39.9*	27.8*
DoGE (Iter 2)	MLDR-13	49.6	44.9	47.0	47.1	72.7	70.5	34.6	58.5	44.4	37.0	73.3	56.7	33.2	25.0
CRISP (n=32/lang)	MLDR-13	57.1	56.7	53.3	52.6	79.5	75.6	44.0	66.6	55.4	48.8	79.3	64.7	36.7	28.9
CRISP (n=32 ² /lang)	MLDR-13	56.6	53.1	53.9	52.0	79.8	75.5	44.9	65.5	54.5	48.6	78.0	62.8	38.5	28.5

Table 11: Average NDCG@10 scores on the MLDR-13 test set across 13 languages. [◊] denotes reproduced results, [†] and [‡] mark the two models compared in the paired significance test, and * indicates a statistically significant difference ($p < 0.05$).

Reptile Update	Train Dataset	Dev Dataset	Fever Test	Fiqa Test	HotpotQA Test
On	BEIR-7	BEIR-1	54.7	33.7	52.9
Off			56.5	28.8	52.7

Table 12: Average NDCG@10 on FEVER, FiQA, and HotpotQA test collection, optimization of scorer is done individually on the development sets without reptile updates.

Update steps l	FEVER	FiQA
0	57.3	29.0
1	59.4	30.9
3	64.9	33.7
5	62.4	33.4
10	62.9	31.2
20	60.6	31.4

Table 13: Effect of the number of update steps l on NDCG@10 when optimizing separately for FEVER and FiQA.

Language	mix1	mix2	mix3	mix4	mix5	Train	Mix Avg.
ar	35.7	35.7	35.9	35.2	34.7	36.2	35.44 $\pm$ 0.5
de	34.2	34.0	33.7	34.4	34.1	34.5	34.08 $\pm$ 0.3
en	29.5	30.2	29.4	30.7	29.3	29.6	29.82 $\pm$ 0.6
es	58.7	59.8	60.7	59.4	59.5	60.3	59.62 $\pm$ 0.7
fr	56.7	57.0	58.4	58.2	56.6	57.2	57.38 $\pm$ 0.9
hi	26.7	25.6	25.0	23.7	23.8	25.1	24.96 $\pm$ 1.3
it	49.5	49.6	48.6	49.5	48.2	49.1	49.08 $\pm$ 0.6
ja	39.4	38.8	38.4	39.8	39.0	38.5	39.08 $\pm$ 0.5
ko	34.7	32.6	32.7	32.9	31.7	32.6	32.92 $\pm$ 1.1
pt	61.3	59.4	61.4	60.8	60.0	61.7	60.58 $\pm$ 0.9
ru	49.8	49.6	50.6	49.8	49.3	51.4	49.82 $\pm$ 0.5
th	26.4	25.0	24.6	25.2	24.3	25.7	25.10 $\pm$ 0.8
zh	17.6	16.8	17.5	17.4	17.8	17.5	17.42 $\pm$ 0.4
Avg.	40.0	39.5	39.8	39.8	39.1	40.0	39.6 $\pm$ 0.3

Table 14: Testing Dev/Test leakage on *bge-m3-retromae* via repeated sampling over the mixture of train and dev splits of MLDR, with evaluation on the MLDR-13 test set. We report NDCG@10 across five mixed-data runs, along with the mix mean and standard deviation.

Dataset	mix1	mix2	mix3	mix4	mix5	Train	Mix Avg.
FEVER	78.1	77.7	78.5	77.5	78.3	76.6	78.02 $\pm$ 0.4
FiQA	30.7	30.6	30.8	30.4	30.8	31.0	30.66 $\pm$ 0.2
HotpotQA	30.4	30.1	30.3	30.4	30.0	30.0	30.24 $\pm$ 0.2
Avg.	46.4	46.1	46.5	46.1	46.4	45.9	46.3 $\pm$ 0.2

Table 15: Testing Dev/Test leakage on *roberta-base* via repeated sampling over the mixture of BEIR train and development splits, with evaluation on the BEIR test collections. We report NDCG@10 across five mixed-data runs, along with the mix mean and standard deviation.




Figure 10: Sampling probability trajectories of MultiDDS and Inf-DDS with varying initialization temperatures during optimization on the FEVER development set. The orange curve denotes the FEVER training set sampling trajectory.




Figure 11: Sampling probability trajectories of MultiDDS and Inf-DDS with varying initialization temperatures during joint optimization on the BEIR-5 development set (DBpedia, FEVER, FiQA, HotpotQA, Quora). The brown curve represents the sampling trajectory for the NFCorpus training set, which is aggressively upsampled by MultiDDS, resulting in degraded overall performance, as shown in Table 9.




Figure 12: Sampling probability trajectories of MultiDDS and Inf-DDS with varying initialization temperatures during optimization on the FiQA development set. The green curve denotes the FiQA training set sampling trajectory. Compared to MultiDDS, which upsamples FiQA due to gradient similarity, Inf-DDS upsamples datasets like MS MARCO and FEVER that are more relevant for performance gains as seen in Table 8.




Figure 13: Sampling probability trajectories of MultiDDS and Inf-DDS with varying initialization temperatures during joint optimization on the HotpotQA development set. The red curve represents the sampling trajectory for the HotpotQA training set, which is being upsampled more by MultiDDS, leading to a better performance as seen in Table 8.




Figure 14: Sampling probability trajectories of MultiDDS and Inf-DDS with varying initialization temperatures during joint optimization on the Dbpedia development set. Since DBpedia is not present in the training set, we see FEVER being upsampled more by Inf-DDS, which should be true since DBpedia and FEVER are very closely related datasets (Thakur et al., 2021).

Sampling Probabilities Over Training Steps




Figure 15: Sampling probability trajectories of Inf-DDS with *Expert* initialization during training of *all-MiniLM-L6-v2* on Sentence Transformers data. Optimization is done jointly on the BEIR-5 development set (DBpedia, FEVER, FiQA, HotpotQA, and Quora).




Figure 16: Sampling probability trajectories of Inf-DDS with $\tau = 1$ initialization during training of *all-MiniLM-L6-v2* on Sentence Transformers data. Optimization is done jointly on the BEIR-5 development set (DBpedia, FEVER, FiQA, HotpotQA, and Quora).




Figure 17: Sampling probability trajectories of Inf-DDS during training of *bge-m3-dense* while jointly optimizing for MLDR-13 dev sets.