

---

# Entropy-Based Adaptive Weighting for Self-Training

---

Xiaoxuan Wang<sup>1</sup> Yihe Deng<sup>1</sup> Derek Mingyu Ma<sup>1</sup> Wei Wang<sup>1</sup>

## Abstract

The mathematical problem-solving capabilities of large language models have become a focal point of research, with growing interests in leveraging self-generated reasoning paths as a promising way to refine and enhance these models. These paths capture step-by-step logical processes while requiring only the correct answer for supervision. The self-training method has been shown to be effective in reasoning tasks while eliminating the need for external models and manual annotations. However, optimizing the use of self-generated data for model training remains an open challenge. In this work, we propose Entropy-Based Adaptive Weighting for Self-Training (EAST), an adaptive weighting strategy designed to prioritize uncertain data during self-training. Specifically, EAST employs a mapping function with a tunable parameter that controls the sharpness of the weighting, assigning higher weights to data where the model exhibits greater uncertainty. This approach guides the model to focus on more informative and challenging examples, thereby enhancing its reasoning ability. We evaluate our approach on GSM8K and MATH benchmarks. Empirical results show that, while the vanilla method yields virtually no improvement (0%) on MATH, EAST achieves around a 1% gain over backbone model. On GSM8K, EAST attains a further 1–2% performance boost compared to the vanilla method. Our codebase is publicly available on GitHub (<https://github.com/mandyyyyyi/east>).

## 1. Introduction

Mathematical reasoning is a key component of Large Language Model (LLM) capabilities, as it directly relates to

---

<sup>1</sup>Department of Computer Science, University of California, Los Angeles. Correspondence to: Xiaoxuan Wang <[xw27@ucla.edu](mailto:xw27@ucla.edu)>.

logical consistency and problem-solving skills (Yu et al., 2023a; Zhang et al., 2024; Gao et al., 2024; Liu et al., 2024). This area has drawn increasing attention because the correctness of a final mathematical answer can provide a direct, verifiable reward signal for reinforcement learning (RL) approaches, enabling LLM-generated reasoning paths for both self-training (Zelikman et al., 2022; Singh et al., 2023; Xiong et al., 2024) and distillation (Ho et al., 2022; Fu et al., 2023; Gou et al., 2023).

The core idea of both self-training and distillation is based on rejection sampling: for each given question, the LLM generates multiple responses and selects the reasoning paths that yield correct answers as positive samples for subsequent fine-tuning (Zelikman et al., 2022; Singh et al., 2023; Luong et al., 2024). Through iterative application of this self-training process, the LLM progressively enhances its performance. Recent studies have explored leveraging negative samples to construct preference pairs for reward models (Hosseini et al., 2024) or directly applying pair-wise alignment methods for fine-tuning (Xu et al., 2024b; Sun et al., 2024; Zhong et al., 2024; Ivison et al., 2024; Saeidi et al., 2024; Xiong et al., 2024).

However, many self-training methodologies treat generated data uniformly, assigning equal importance to all generated examples. Such approaches may overlook the varying educational value of different data points, which can potentially impede the model’s ability to prioritize the most informative data and possibly limits its overall learning effectiveness. This observation raises a question: could reweighting training data during self-training improve reasoning capabilities? If so, which data should be prioritized, and to what extent should it be emphasized?

In the self-training pipeline for reasoning tasks, additional training on already well-understood questions brings minimal gains and risks overfitting the model to simpler data. Instead, focusing on challenging questions—where the model struggles—promises more efficient learning (Huang et al., 2022; Singh et al., 2023). Moreover, large language models can exhibit resistance to updating their predictions, particularly in cases where they demonstrate high confidence. In contrast, guiding the model to focus on areas of uncertainty enhances its training effectiveness (Kumar et al., 2024; Li et al., 2024). To address this gap, we introduce Entropy-




Figure 1. Comparison between the traditional self-training pipeline and EAST. The LLM generates  $n$  responses per question, clustered by final answers. Questions with all incorrect answers are discarded. Self-training fine-tunes uniformly on the rest, while EAST assigns higher weights to questions with diverse (uncertain) answers and lower weights to consistent (confident) ones.

Based Adaptive Weighting for Self-Training (EAST), a novel method that assigns adaptive weights to training data during self-training based on model uncertainty, measured via the entropy of the model’s sample distribution for a given question. Specifically, given multiple samples generated by an LLM for a question, EAST clusters these samples by their final answers and computes the entropy over the resulting cluster-based distribution. EAST then applies a mapping function that transforms the entropy value into a bounded weight under predefined constraints. This function includes a tunable parameter that controls the sharpness of the weighting, allowing flexible emphasis on uncertain data. By assigning higher weights to high-entropy data—those reflecting greater model uncertainty—EAST encourages the model to focus on more informative and challenging examples during training. Prioritizing such examples not only enhances reasoning capability but also helps prevent overfitting to overconfident data. Moreover, EAST is a flexible framework that supports both iterative self-training and integration with various loss functions, making it broadly applicable across different training settings.

We evaluate EAST by incorporating it into SFT, DPO (Rafailov et al., 2024), and KTO (Ethayarajh et al., 2024) loss functions on two mathematical reasoning benchmarks GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). EAST achieves notable performance gains of 5.6% on GSM8K and approximately 1% on MATH over the default backbone model, substantially outperforming vanilla SFT, which yields only a 3.9% improvement on GSM8K and no gain on MATH. A similar trend is observed for both DPO and KTO, with performance improvements of up to 1.7% on GSM8K and 2.1% on MATH compared to the vanilla methods. We further show that EAST consistently surpasses vanilla method through iterative training. In addition, we demonstrate the effectiveness of entropy-based weighting, which outperforms other weighting strategies by better leveraging uncertain data and reducing reliance

on overconfident data during training, thereby enhancing reasoning capabilities.

Our contributions are summarized as follows:

- **Entropy-Based weighting:** a new weighting strategy that leverages uncertainty information, derived from the entropy of the model’s sample distribution over the training data
- **Mapping function:** a novel mapping function that controls the extent to which higher uncertain data are weighted
- **Experimental evaluation :** EAST further boosts self-training performance compared to the vanilla method.

## 2. Preliminaries

We consider a large language model (LLM) parameterized by  $\theta$ , denoted  $p_\theta$ . Given a prompt  $x = [x_1, \dots, x_n]$ , the model generates a response  $y = [y_1, \dots, y_m]$  via an autoregressive factorization:

$$p_\theta(y | x) = \prod_{j=1}^m p_\theta(y_j | x, y_{<j}),$$

where  $y_{<j} = [y_1, \dots, y_{j-1}]$ .

**Self-Training Pipeline.** Self-training addresses the scarcity of human-annotated data by leveraging the target model to generate completion paths (Zelikman et al., 2022; Singh et al., 2023; Chen et al., 2024). Formally, under mathematical context, given a dataset of input-output pairs  $((x_i, y_i)_{i=1}^N)$ , where  $(x_i)$  represents a mathematical question and  $(y_i)$  is its corresponding ground truth answer, we aim to introduce an intermediate reasoning path  $(r_i)$  that delineates the logical steps from  $(x_i)$  to  $(y_i)$ . Let  $(p_\theta(r_i | x_i))$  denote the model’s distribution over possible reasoning paths, parameterized by  $(\theta)$ . The self-training process involves:

1. Sampling reasoning paths ( $\hat{r}_i \sim p_\theta(\cdot | x_i)$ )
2. Evaluating the correctness of ( $\hat{r}_i$ ) by verifying if it leads to the ground truth ( $y_i$ )
3. Updating the training set with validated triples  $((x_i, y_i, \hat{r}_i)_{i=1}^M)$
4. Updating model parameters ( $\theta$ ) through iterative training

For supervised fine-tuning (SFT), only sample paths that yield correct answers are incorporated into the training data. Alignment methods such as direct preference optimization (DPO, Rafailov et al., 2024) utilize both correct and incorrect sample paths to learn from contrastive preferences, where correct paths serve as positive pairs and incorrect ones as negative pairs.

### 3. Method

In this section, we introduce EAST, a novel weighting method that prioritizes uncertain data within the self-training pipeline. We begin by presenting the entropy-based weighting strategy, followed by the proposed mapping function, and conclude with the final loss objective. Figure 1 demonstrates the comparison between the traditional self-training pipeline and EAST. Figure 2 represents the detailed framework of EAST.

#### 3.1. Entropy-Based Weight

Many studies have found that large language models (LLMs) tend to be resistant to changing their predictions, particularly when they are highly confident in their responses (Kumar et al., 2024; Yang et al., 2024c; Li et al., 2024). Therefore, guiding the model to focus on areas where it lacks confidence or is uncertain becomes a natural next step. Research indicates that prioritizing learning from uncertain questions—rather than those where the model is stubborn—leads to improved reasoning capabilities. (Kumar et al., 2024; Li et al., 2024).

Based on this observation, we introduce an entropy-based weighting approach that encourages models to focus on learning from uncertain data. The key insight is that questions with higher entropy reflect greater uncertainty in the model’s predictions, indicating a lack of strong preference among possible answers. By prioritizing high-entropy questions during training, the model is encouraged to focus on informative and challenging examples, which enhances reasoning capabilities and helps prevent overfitting to overconfident examples. We further demonstrate the advantage of entropy-based weighting over alternative weighting strategies in Section 4.3, including accuracy-based weighting, which considers the proportion of correct answers (accuracy

ratio), and rejection-based weighting, which captures the dominance of the most frequent incorrect answer (dominant incorrect ratio).

In the self-training pipeline, we generate  $n$  samples for each question and cluster them based on their final answers, with each cluster representing a distinct answer. The number of clusters for a given question depends on the diversity of the model’s outputs, which we denote as  $k_i$  (where  $k_i \leq n$ ) for question  $x_i$ . The underlying assumption is that samples leading to the same final answer tend to share similar reasoning patterns. Thus, each cluster reflects the model’s implicit preference for a particular reasoning path. A larger number of sparse clusters ( $k_i$ ) indicates greater model uncertainty for that question, while a distribution concentrated in a single cluster suggests higher model confidence. The model’s uncertainty for a given question  $x_i$  is quantified through the entropy value, computed over the  $k_i$  answer clusters as follows:

$$H(x_i) = - \sum_{j=1}^{k_i} p_j \log p_j \quad (1)$$

where  $p_j$  denotes the proportion of samples in cluster  $j$  relative to the total number of samples for  $x_i$ . For simplicity, we denote  $h_i = H(x_i)$ .

#### 3.2. Mapping Function

Given the entropy value  $h_i$  on each question, our goal is to map this entropy value to a weight applied to the model loss using a function  $f$ . This mapping function  $f$  must satisfy two constraints: (1) *Non-negativity*—all transformed weights must be non-negative to ensure proper model training; (2) *Normalization*—the transformed weights should have an average of 1 to prevent unintended effects on the learning rate, formally:

$$\min(f(h)) \geq 0, \quad \frac{1}{N} \sum_{i=1}^N f(h_i) = 1. \quad (2)$$

**Attempt 1.** A straightforward mapping function is the mean-division function:

$$f(h) = \frac{h}{\mu}, \quad \mu = \frac{1}{N} \sum_{i=1}^N h_i \quad (3)$$

which satisfies the normalization constraint. However, this approach lacks tunable parameters to control the distribution of transformed values.

**Attempt 2.** To allow for control over the distribution of transformed values, we introduce a new parameter  $R = f(\max(h)) - f(\min(h))$  that represents the range of mapped values. Therefore, instead of applying a fixed compression ratio (as in the mean-division function) to entropy values, we allow a controllable compression ratio  $a$

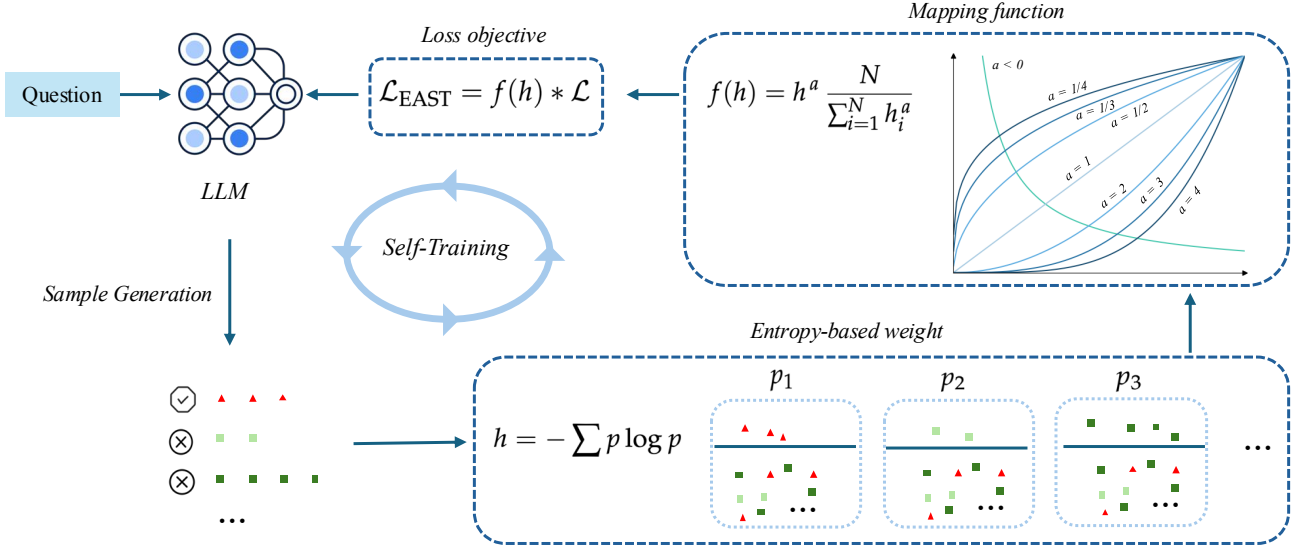


Figure 2. The framework of EAST. For each training question, the LLM generates  $n$  responses, clustered by final answers. Entropy value is computed from the cluster distribution, transformed via mapping function, and integrated as weight into the loss objective.

that adapts the new output range  $R$  and the original range  $(\max(h) - \min(h))$ :

$$f(h) = ah + b, \quad a = \frac{R}{\max(h) - \min(h)}, \quad b = 1 - a\mu. \quad (4)$$

Here,  $b$  is determined by the normalization constraint:  $\frac{1}{N} \sum_{i=1}^N (ah_i + b) = a\mu + b = 1$ , which gives  $b = 1 - a\mu$ . The non-negativity constraint ( $\min(f(h)) \geq 0$ ) requires  $a(\min(h) - \mu) + 1 \geq 0$ , yielding  $R \leq \frac{\max(h) - \min(h)}{\mu - \min(h)}$ .

While this linear approach offers some control, it has two key limitations: (1) the output range  $R$  is upper-bounded by the non-negativity requirement, and (2) the linear mapping does not allow for "curvature" control to amplify or compress differences between entropy values.

**Attempt 3 (Final).** To address the non-negativity constraint, we propose mapping the transformed values into the exponential space by applying a logarithmic transformation to the entropy values:

$$f(h) = e^{a \ln h + b} = h^a \cdot e^b, \quad (5)$$

where exponent parameter  $a$  controls the curvature of the transformation, providing flexibility in how the entropy values are reshaped. This formulation automatically ensures  $f(h)$  is non-negative for  $h > 0$ . Substituting into the normalization constraint and solving for  $b$ :

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N f(h_i) &= \frac{1}{N} \sum_{i=1}^N h_i^a e^b = e^b \frac{1}{N} \sum_{i=1}^N h_i^a = 1 \\ \Rightarrow e^b &= \frac{N}{\sum_{i=1}^N h_i^a} \Rightarrow b = \ln \left( \frac{N}{\sum_{i=1}^N h_i^a} \right) \end{aligned}$$

Therefore, our final mapping function is:

$$f(h) = h^a \frac{N}{\sum_{i=1}^N h_i^a}. \quad (6)$$

The exponent parameter  $a$  provides curvature control over the transformation:  $f(h)$  enhances differences between weights when  $a > 1$ ;  $f(h)$  compresses differences when  $0 < a < 1$ ;  $f(h)$  inverts the weight distribution when  $a < 0$ .

---

**Algorithm 1** Entropy-Based Adaptive Weighting for Self-Training (EAST)

---

**Input:** Initial model parameters  $\theta$ , training data  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ , exponent parameter  $a$ , maximum iterations  $T$

**Output:** Trained model parameters  $\hat{\theta}$

**for**  $t = 1$  **to**  $T$  **do**

**foreach** question  $x_i \in \mathcal{D}$  **do**

        Generate  $n$  responses and cluster them into  $k_i$  groups by final answers with proportions  $p_1, \dots, p_{k_i}$  per group. Compute entropy:  $h_i = -\sum_{j=1}^{k_i} p_j \log p_j$

    Compute coefficient:  $e^b = \frac{N}{\sum_{i=1}^N h_i^a}$

**foreach** question  $x_i \in \mathcal{D}$  **do**

        Compute weight:  $f(h_i) = h_i^a \cdot e^b$

        Update loss function:  $\mathcal{L}_{EAST}(\theta; x_i) = f(h_i) \cdot \mathcal{L}(\theta; x_i)$

    Train model by minimizing  $\mathcal{L}_{EAST}$

**return**  $\hat{\theta}$

---



### 3.3. Loss Objective

The resulting weight is then integrated into the loss objective as:

$$\mathcal{L}_{\text{EAST}}(\theta) = f(h) \cdot \mathcal{L}(\theta), \quad \text{where } h = H(x) \quad (7)$$

where  $f(h)$  is the mapping function applied to the entropy value  $H(x)$ , and  $\mathcal{L}(\theta)$  denotes the base loss. EAST is flexible and can be seamlessly applied to various loss functions (e.g., SFT or DPO). Furthermore, it naturally supports iterative training by repeating the weighting and fine-tuning process. The full procedure is detailed in Algorithm 1.

## 4. Experiment

In this section, we present the experiment setup in Section 4.1 and main results in Section 4.2. Then, we provide further ablation study in Section 4.3.

### 4.1. Experiment Setup

**Dataset.** We evaluate EAST on two mathematical benchmarks: MATH (Hendrycks et al., 2021) and GSM8K (Cobbe et al., 2021). For training data, we prompt the backbone model to generate 128 samples per question and randomly select a positive-negative pair based on answer correctness, with the positive drawn from correct answers and the negative from incorrect ones. For evaluation, we note minor performance variations with vLLM across GPU types. For reproducibility and fair comparison, all results use the same GPU with temperature 0. We adapt the evaluation pipeline of Yang et al. (2024a).

**Baseline.** We evaluate EAST across three loss functions: SFT, DPO (Rafailov et al., 2024), and KTO (Ethayarajh et al., 2024), which correspond to learning from positive samples, paired samples, and unpaired samples. In addition to the vanilla method, we incorporate weighting baselines that capture local uncertainty information. Specifically, local uncertainty information refers to model uncertainty derived exclusively from the token-level probabilities of a single selected sample response. This metric captures the uncertainty within an individual response, without accounting for the full distribution of all generated responses for a given question. One baseline uses the perplexity score for local information weighting (denoted as **LW(P)**), which is normalized within each batch to ensure stability and fair comparison. Another baseline (denoted as **LW(W)**) is inspired by WPO (Zhou et al., 2024), which computes adaptive weights based on the log-likelihood of the sample response. Detailed formulations for both baselines are provided in the Appendix B.1.

**Model Configuration.** We conduct experiments systematically based on two backbone models: Llama-3.2-1B-Instruct and

Llama-3.1-8B-Instruct. For SFT, we use a learning rate of 2e-6 for 1B on GSM8K and MATH datasets. We adapt LoRA for 8B model with learning rate as 5e-5 for GSM8K and 2e-5 for MATH. For DPO, we adapt a learning rate of 2e-7 for 1B with  $\beta = 0.01$  and 2e-6 for 8B using LoRA with  $\beta = 0.1$  for both datasets. For KTO, we adapt a learning rate of 2e-7 for 1B with  $\beta = 0.05$  and 2e-6 for 8B using LoRA with  $\beta = 0.1$  for both datasets. For both baselines and EAST, we use the same set of hyperparameters as the vanilla method to ensure fair comparison. Each model is trained for three epochs with a batch size of 16 and warmup ratio of 0.1. For Llama-3.1-8B-Instruct, we apply LoRA with a rank of 16 and a LoRA alpha of 16. All models are trained using bfloat16 precision, and we use the AdamW optimizer. We adapt the exponent parameters  $a$  in the range  $[-3, 3]$  to fully investigate the functionality of the mapping function. Detailed hyperparameter study is provided in Appendix B.2.

### 4.2. Experiment Results

Table 1 presents the performance in terms of accuracy score of EAST compared to the vanilla method and baselines on the GSM8K and MATH benchmarks during the first iteration. The impact of different exponent parameters  $a$  is shown in Figure 3. Additionally, experiment results from iterative learning are presented in Figure 4. We have the following observations:

**Observation 1: EAST outperforms the vanilla method and baselines.** As shown in Table 1, EAST consistently improves performance compared to the vanilla method and baselines across both benchmarks under all loss functions. The vanilla SFT method struggles to outperform the default backbone model when trained on self-generated data, especially on the challenging MATH dataset. For instance, SFT achieves accuracies of 28.4% and 50.0% on MATH using LLaMA-3.2-1B and LLaMA-3.1-8B, respectively, which are slightly lower than the corresponding default model performances of 28.5% and 50.4%. In comparison, EAST improves the results to 29.4% and 51.2%, demonstrating over a 1% absolute gain relative to SFT by focusing on more informative training examples. A similar trend is also observed on KTO: the vanilla KTO achieves only 83.9% and 48.9% on the 8B model, while EAST boosts the performance to 85.1% and 51.0%, representing gains of 1.2% and 2.1%, respectively. Integrating EAST with the DPO loss function also leads to consistent gains. For example, on GSM8K with the LLaMA-3.2-1B model, EAST improves performance from 50.2% to 51.9%.

We also evaluate baselines (LW(W) and LW(P)) that use local uncertainty information of model. These approaches rely on the likelihood of next-token prediction for a given

Table 1. Experimental results in terms of accuracy(%) on GSM8K and MATH benchmarks. The best performance under each loss category is highlighted in **bold**. Significant boosts ( $\geq 1\%$ ) of EAST over both the vanilla method and baselines are underlined.

Setting	LLaMA-3.2-1B			LLaMA-3.1-8B		
	GSM8K(%)	MATH(%)	AVG(%)	GSM8K(%)	MATH(%)	AVG(%)
<i>default</i>	46.2	28.5	37.3	82.8	50.4	66.6
SFT	50.1	28.4	39.2	85.0	50.0	67.5
+LW(W)	50.9	28.5	39.7	84.8	50.9	67.8
+LW(P)	51.2	28.4	39.8	85.1	50.8	68.0
+EAST	<b>51.8</b>	<b>29.4</b>	<b>40.6</b>	<b>86.1</b>	<b>51.2</b>	<b>68.6</b>
DPO	50.2	28.7	39.5	84.6	50.1	67.5
+LW(W)	50.9	28.1	39.5	85.1	50.2	67.6
+LW(P)	50.3	28.4	39.4	85.2	50.8	68.0
+EAST	<u>51.9</u>	<u>29.7</u>	<u>40.8</u>	<b>85.4</b>	<b>50.9</b>	<b>68.1</b>
KTO	<b>53.0</b>	28.8	40.9	83.9	48.9	66.4
+LW(W)	52.9	28.2	40.6	84.3	49.1	66.7
+LW(P)	52.9	28.9	40.9	83.9	49.1	66.5
+EAST	<b>53.0</b>	<u>29.9</u>	<b>41.5</b>	<b>85.1</b>	<u>51.0</u>	<u>68.1</u>

sample response, rather than the overall sample distribution, which also accounts for the correctness of the sample path. Results show that EAST outperforms both local information baselines across all loss functions and benchmarks. Notably, on the MATH dataset using the LLaMA-3.2-1B model, local weighting baselines achieve only 28.1% (LW(W)) and 28.4% (LW(P)), which is lower than both the vanilla DPO (28.7%) and the default model (28.5%). In contrast, EAST achieves a significantly higher score of 29.7%, suggesting that local weighting may be more sensitive to token-level noise and potentially limiting its training effectiveness.

**Observation 2: Weighting more on uncertain data contributes to performance improvement.** Figure 3 demonstrates the accuracy score of different parameters  $a$  using SFT method on both benchmarks using LLaMA-3.2-1B model. The figure demonstrates that while performance varies with different values of  $a$ , the best results are achieved when  $a > 0$  for both datasets. The model achieves peak accuracy on the GSM8K dataset at  $a = 1.5$  with 51.8%, compared to 50.6% and 50.8% when  $a = -0.75$  and  $a = -1$ , respectively. Similarly, for the MATH dataset, performance reaches 29.3% at  $a = 1$  and 29.4% at  $a = 3$ , outperforming the 28.1% and 28.5% observed at  $a = -1$  and  $a = -0.75$ , respectively. These results suggest that prioritizing uncertain data helps the model enhance its reasoning ability during training, leading to improved performance.

**Observation 3: EAST demonstrates consistent benefits in iterative training.** To further investigate the performance

of EAST in iterative learning, we conduct experiments with iterations  $T = 3$  using LLaMA-3.2-1B on both the MATH and GSM8K datasets, with results presented in Figure 4. EAST consistently outperforms vanilla SFT across iterations on both datasets. Notably, EAST maintains strong performance over time, while vanilla SFT appears to overfit on self-generated data in GSM8K. Although both methods struggle with iterative learning on the MATH dataset, EAST still demonstrates a relative advantage.

#### 4.3. Ablation Study: Effect of Accuracy-Based and Reject-Based Weights

To further investigate the effectiveness of entropy-based weighting, we compare it to alternative weighting strategies grounded in other distributional metrics. Noting that entropy is typically low when a single answer—whether correct or incorrect—dominates the distribution, we explore two complementary approaches: *accuracy-based weighting*, which considers the proportion of the correct answer, and *rejection-based weighting*, which measures the dominance of the most frequent incorrect answer.

##### 4.3.1. ACCURACY-BASED WEIGHTS.

Accuracy-based weighting leverages the accuracy ratio of model for each question to determine the corresponding weight. Specifically, in the self-training pipeline with  $n$  samples for each question in the training data and the accuracy score is computed based on the proportion of correct




Figure 3. Performance(accuracy (%)) of various exponent parameters  $a$  on GSM8K and MATH datasets using LLaMA-3.2-1B.




Figure 4. Comparison of iterative learning performance (accuracy (%)) between vanilla SFT and EAST on LLaMA-3.2-1B.

predictions:  $A(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(y_i = y^*)$ , where  $y_i$  represents the  $i$ -th sampled prediction and  $y^*$  denotes the ground truth. For notational simplicity, let  $s_i = 1 - A(x_i)$  represent the inverse accuracy score for question  $i$  and the weight is aggregated using mapping function  $f(s_i)$ , as detailed in Section 3.2. Intuitively, when  $s_i$  is large, the model faces challenges when solving the problems.

#### 4.3.2. REJECTED-BASED WEIGHT.

Recent studies indicate that large language models (LLMs) struggle with self-correction, particularly when they generate responses with high confidence (Kumar et al., 2024; Yang et al., 2024c). To further investigate this phenomenon, we propose a novel weighting scheme that prioritizes the most “stubborn” questions—those for which the model repeatedly produces the same incorrect answers. Specifically, for all samples that yield incorrect answers, we partition them into  $k$  clusters, where each cluster corresponds to a distinct final answer. Next, we calculate the proportion  $p$  of each incorrect answer cluster and identify the most frequent (dominant) mistake:  $R(x) = \max_{j \in [k]} p_j$ . For notational simplicity, let  $r_i = R(x_i)$  represent the inverse accuracy score for question  $i$  and the weight is aggregated using mapping function  $f(r_i)$ , as detailed in Section 3.2.

#### 4.3.3. EXPERIMENT RESULT AND ANALYSIS.

As shown in Equation 1, both  $A(x_i)$  and  $R(x_i)$  can be interpreted as components of the probability distribution  $H(x_i)$  over predicted answers. According to the equation, when either  $A(x_i)$  or  $R(x_i)$  is large, the entropy  $H(x_i)$  tends to be low, reflecting greater certainty in the model’s

predictions. This relationship is further illustrated in Figure 5, where higher entropy values are associated with larger  $(1 - A(x_i))$  (i.e., lower accuracy) and smaller  $R(x_i)$  values (i.e., greater diversity in incorrect predictions). However, when  $(1 - A(x_i))$  becomes large—exceeding 0.8, for instance—this does not necessarily imply a low  $R(x_i)$ . In fact, Figure 5 shows that many of these low-accuracy samples still exhibit high  $R(x_i)$  values, indicating repeated, confident errors. As a result, applying the mapping function  $f$  to such cases may overemphasize these “stubborn” questions during training, potentially skewing the learning dynamics and degrading overall performance. In contrast, entropy-based weighting effectively addresses this problem by automatically assigning lower weights to cases where a single incorrect answer dominates.

Empirical results of the three weighting strategies on both datasets are reported in Figure 5, using SFT and LLaMA-3.2-1B. The results show that entropy-based weighting outperforms other strategies on both datasets. In contrast, reject-based weighting consistently yields the lowest performance across both benchmarks, while accuracy-based weighting achieves comparable results but exhibits certain limitations.

## 5. Related Work

### 5.1. Self-Training on Mathematical Reasoning.

Mathematical reasoning has emerged as a critical evaluation benchmark for Large Language Models (LLMs), as it directly correlates with logical reasoning capabilities and provides clear assessment metrics (Azerbayev et al., 2023; Wang et al., 2023; Zhang et al., 2024; Gao et al., 2024; Liu

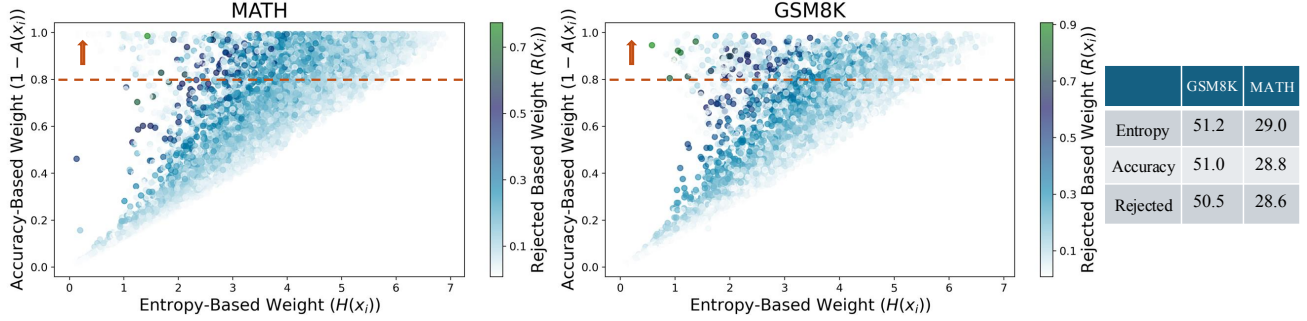


Figure 5. The figure illustrates the distribution of training data in entropy-based, accuracy-based, and rejected-based values. Each point represents a training example  $(x_i)$ , with coordinates  $(H(x_i), 1 - A(x_i))$  for entropy-based and accuracy-based values, and color indicating the rejected-based value  $(R(x_i))$ . The accompanying table reports the performance (accuracy(%)) of three weighting strategies on the GSM8K and MATH datasets.

et al., 2024). Traditional methods rely on carefully curated manual datasets as demonstrations for fine-tuning (Yue et al., 2023; Yu et al., 2023a; Luo et al., 2023). As high-quality annotated data are expensive, numerous studies leverage rephrasing methods to augment datasets (Deng et al., 2023; Yu et al., 2023a), or employ strong LLMs to generate synthetic data for knowledge distillation (Taori et al., 2023; Chiang et al., 2023; Ho et al., 2022; Fu et al., 2023; Gou et al., 2023).

Recently, several studies have explored using the target model to generate training data and enhance its performance through self-training (Zelikman et al., 2022; Singh et al., 2023; Hosseini et al., 2024; Yu et al., 2023b; Chen et al., 2024; Kumar et al., 2024; Tao et al., 2024), while others extend such techniques for pair-wise alignment methods by leveraging negative samples generated from previous iterations (Tajwar et al., 2024; Xu et al., 2024b; Sun et al., 2024; Zhong et al., 2024; Ivison et al., 2024; Xiong et al., 2024; Xie et al., 2024; Pang et al., 2025). For instance, Sun et al. (2024) compare REST-EM with iterative DPO in a self-training pipeline, and Xiong et al. (2024) employs the multi-turn reasoning path for iterative learning process. For this work, we further optimize self-generated data usage by incorporating weighting strategies to improve reasoning capabilities.

## 5.2. Alignment Method.

Reinforcement Learning from Human Feedback (RLHF) has emerged as an essential framework for aligning machine learning models with human preferences, emphasizing the importance of post-training optimization. Direct Preference Optimization (DPO) is a widely recognized method for alignment (Rafailov et al., 2024). Recent studies have explored derivation of DPO (Xu et al., 2024a; Meng et al., 2024; Azar et al., 2024). For example, KTO directly maximizes the utility of generated outputs instead of focus-

ing on the log-likelihood of preferences (Ethayarajh et al., 2024). Some other studies focused on applying local weighting (Zhou et al., 2024) or reward weighting upon DPO (Adler et al., 2024; Xiao et al., 2024; Yang et al., 2024b). RPO incorporates reward gaps into preference learning to mitigate overfitting and better capture nuanced response quality (Adler et al., 2024). However, it relies on an external reward model to assign weights to preference pairs. In contrast, WPO reweights preference pairs based on their likelihood under the current policy (Zhou et al., 2024). Nevertheless, WPO relies solely on local information of the given sample response without considering the overall sample distribution or controlling weight distribution skewness.

## 6. Conclusion

This paper introduces EAST, an Entropy-based Adaptive Self-Training method that dynamically assigns weights to training samples based on their uncertainty, with the goal of enhancing the reasoning capabilities of large language models (LLMs) during the self-training process. The core idea behind EAST is to leverage entropy—a measure of prediction uncertainty—to prioritize more ambiguous or challenging examples, which are often more informative for learning. To achieve this, EAST applies a tunable power-law mapping function that transforms the entropy values into sample weights, enabling fine-grained control over the contribution of each sample during model optimization. We conduct extensive experiments on two widely-used benchmarks for mathematical reasoning—GSM8K and MATH—to evaluate the effectiveness of EAST. The results demonstrate consistent improvements over standard self-training baselines across multiple metrics, validating both the robustness and generality of the approach. Notably, EAST is compatible with a wide range of self-training pipelines and can be integrated as a lightweight yet powerful plug-in module.



## References

- Adler, B., Agarwal, N., Aithal, A., Anh, D. H., Bhattacharya, P., Brundyn, A., Casper, J., Catanzaro, B., Clay, S., Cohen, J., et al. Nemotron-4 340b technical report. *arXiv preprint arXiv:2406.11704*, 2024.
- Azar, M. G., Guo, Z. D., Piot, B., Munos, R., Rowland, M., Valko, M., and Calandriello, D. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pp. 4447–4455. PMLR, 2024.
- Azerbaiyev, Z., Schoelkopf, H., Paster, K., Santos, M. D., McAleer, S., Jiang, A. Q., Deng, J., Biderman, S., and Welleck, S. Llemma: An open language model for mathematics. *arXiv preprint arXiv:2310.10631*, 2023.
- Chen, Z., Deng, Y., Yuan, H., Ji, K., and Gu, Q. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*, 2024.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6, 2023.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Deng, Y., Zhang, W., Chen, Z., and Gu, Q. Rephrase and respond: Let large language models ask better questions for themselves. *arXiv preprint arXiv:2311.04205*, 2023.
- Ethayarajh, K., Xu, W., Muennighoff, N., Jurafsky, D., and Kiela, D. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.
- Fu, Y., Peng, H., Ou, L., Sabharwal, A., and Khot, T. Specializing smaller language models towards multi-step reasoning. In *International Conference on Machine Learning*, pp. 10421–10430. PMLR, 2023.
- Gao, B., Song, F., Yang, Z., Cai, Z., Miao, Y., Dong, Q., Li, L., Ma, C., Chen, L., Xu, R., et al. Omni-math: A universal olympiad level mathematic benchmark for large language models. *arXiv preprint arXiv:2410.07985*, 2024.
- Gou, Z., Shao, Z., Gong, Y., Yang, Y., Huang, M., Duan, N., Chen, W., et al. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint arXiv:2309.17452*, 2023.
- Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Song, D., and Steinhardt, J. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.
- Ho, N., Schmid, L., and Yun, S.-Y. Large language models are reasoning teachers. *arXiv preprint arXiv:2212.10071*, 2022.
- Hosseini, A., Yuan, X., Malkin, N., Courville, A., Sordoni, A., and Agarwal, R. V-star: Training verifiers for self-taught reasoners. *arXiv preprint arXiv:2402.06457*, 2024.
- Huang, J., Gu, S. S., Hou, L., Wu, Y., Wang, X., Yu, H., and Han, J. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*, 2022.
- Iverson, H., Wang, Y., Liu, J., Wu, Z., Pyatkin, V., Lambert, N., Smith, N. A., Choi, Y., and Hajishirzi, H. Unpacking dpo and ppo: Disentangling best practices for learning from preference feedback. *arXiv preprint arXiv:2406.09279*, 2024.
- Kumar, A., Zhuang, V., Agarwal, R., Su, Y., Co-Reyes, J. D., Singh, A., Baumli, K., Iqbal, S., Bishop, C., Roelofs, R., et al. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.
- Li, L., Chen, Z., Chen, G., Zhang, Y., Su, Y., Xing, E., and Zhang, K. Confidence matters: Revisiting intrinsic self-correction capabilities of large language models. *arXiv preprint arXiv:2402.12563*, 2024.
- Liu, H., Zheng, Z., Qiao, Y., Duan, H., Fei, Z., Zhou, F., Zhang, W., Zhang, S., Lin, D., and Chen, K. Math-bench: Evaluating the theory and application proficiency of llms with a hierarchical mathematics benchmark. *arXiv preprint arXiv:2405.12209*, 2024.
- Luo, H., Sun, Q., Xu, C., Zhao, P., Lou, J., Tao, C., Geng, X., Lin, Q., Chen, S., and Zhang, D. Wizard-math: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*, 2023.
- Luong, T. Q., Zhang, X., Jie, Z., Sun, P., Jin, X., and Li, H. Reft: Reasoning with reinforced fine-tuning. *arXiv preprint arXiv:2401.08967*, 2024.
- Meng, Y., Xia, M., and Chen, D. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*, 2024.
- Pang, R. Y., Yuan, W., He, H., Cho, K., Sukhbaatar, S., and Weston, J. Iterative reasoning preference optimization. *Advances in Neural Information Processing Systems*, 37: 116617–116637, 2025.

- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.
- Saeidi, A., Verma, S., and Baral, C. Insights into alignment: Evaluating dpo and its variants across multiple tasks. *arXiv preprint arXiv:2404.14723*, 2024.
- Singh, A., Co-Reyes, J. D., Agarwal, R., Anand, A., Patil, P., Liu, P. J., Harrison, J., Lee, J., Xu, K., Parisi, A., et al. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*, 2023.
- Sun, Z., Yu, L., Shen, Y., Liu, W., Yang, Y., Welleck, S., and Gan, C. Easy-to-hard generalization: Scalable alignment beyond human supervision. *arXiv preprint arXiv:2403.09472*, 2024.
- Tajwar, F., Singh, A., Sharma, A., Rafailov, R., Schneider, J., Xie, T., Ermon, S., Finn, C., and Kumar, A. Preference fine-tuning of llms should leverage suboptimal, on-policy data. *arXiv preprint arXiv:2404.14367*, 2024.
- Tao, Z., Lin, T.-E., Chen, X., Li, H., Wu, Y., Li, Y., Jin, Z., Huang, F., Tao, D., and Zhou, J. A survey on self-evolution of large language models. *arXiv preprint arXiv:2404.14387*, 2024.
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., and Hashimoto, T. B. Stanford alpaca: An instruction-following llama model, 2023.
- Wang, X., Hu, Z., Lu, P., Zhu, Y., Zhang, J., Subramaniam, S., Loomba, A. R., Zhang, S., Sun, Y., and Wang, W. Scibench: Evaluating college-level scientific problem-solving abilities of large language models. *arXiv preprint arXiv:2307.10635*, 2023.
- Xiao, W., Wang, Z., Gan, L., Zhao, S., He, W., Tuan, L. A., Chen, L., Jiang, H., Zhao, Z., and Wu, F. A comprehensive survey of datasets, theories, variants, and applications in direct preference optimization. *arXiv preprint arXiv:2410.15595*, 2024.
- Xie, C., Huang, Y., Zhang, C., Yu, D., Chen, X., Lin, B. Y., Li, B., Ghazi, B., and Kumar, R. On memorization of large language models in logical reasoning. *arXiv preprint arXiv:2410.23123*, 2024.
- Xiong, W., Dong, H., Ye, C., Wang, Z., Zhong, H., Ji, H., Jiang, N., and Zhang, T. Iterative preference learning from human feedback: Bridging theory and practice for rlhf under kl-constraint. In *Forty-first International Conference on Machine Learning*, 2024.
- Xu, H., Sharaf, A., Chen, Y., Tan, W., Shen, L., Van Durme, B., Murray, K., and Kim, Y. J. Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation. *arXiv preprint arXiv:2401.08417*, 2024a.
- Xu, S., Fu, W., Gao, J., Ye, W., Liu, W., Mei, Z., Wang, G., Yu, C., and Wu, Y. Is dpo superior to ppo for llm alignment? a comprehensive study. *arXiv preprint arXiv:2404.10719*, 2024b.
- Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024a.
- Yang, Z., Wan, F., Zhong, L., Shi, T., and Quan, X. Weighted-reward preference optimization for implicit model fusion. *arXiv preprint arXiv:2412.03187*, 2024b.
- Yang, Z., Zhang, Y., Wang, Y., Xu, Z., Lin, J., and Sui, Z. Confidence vs critique: A decomposition of self-correction capability for llms. *arXiv preprint arXiv:2412.19513*, 2024c.
- Yu, L., Jiang, W., Shi, H., Yu, J., Liu, Z., Zhang, Y., Kwok, J. T., Li, Z., Weller, A., and Liu, W. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023a.
- Yu, X., Peng, B., Galley, M., Gao, J., and Yu, Z. Teaching language models to self-improve through interactive demonstrations. *arXiv preprint arXiv:2310.13522*, 2023b.
- Yue, X., Qu, X., Zhang, G., Fu, Y., Huang, W., Sun, H., Su, Y., and Chen, W. Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint arXiv:2309.05653*, 2023.
- Zelikman, E., Wu, Y., Mu, J., and Goodman, N. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.
- Zhang, R., Jiang, D., Zhang, Y., Lin, H., Guo, Z., Qiu, P., Zhou, A., Lu, P., Chang, K.-W., Gao, P., et al. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? *arXiv preprint arXiv:2403.14624*, 2024.
- Zhong, H., Feng, G., Xiong, W., Zhao, L., He, D., Bian, J., and Wang, L. Dpo meets ppo: Reinforced token optimization for rlhf. *arXiv preprint arXiv:2404.18922*, 2024.
- Zhou, W., Agrawal, R., Zhang, S., Indurthi, S. R., Zhao, S., Song, K., Xu, S., and Zhu, C. Wpo: Enhancing rlhf with weighted preference optimization. *arXiv preprint arXiv:2406.11827*, 2024.

## Supplementary Material for EAST

### A. Reproducibility

All code will be publicly available in the GitHub. All results are evaluated using the NVIDIA RTX A6000 GPU, following the evaluation pipeline of Yang et al. (2024a).

### B. Experiment Setup

#### B.1. Baseline

For local uncertainty information weighting, we use the standard perplexity(LW(P)):

$$\text{PPL}(x, y) = \exp \left( -\frac{1}{|y|} \sum_{t=1}^{|y|} \log \pi_{\theta}(y_t \mid x, y_{<t}) \right), \quad (8)$$

where  $\pi_{\theta}(y_t \mid x, y_{<t})$  denotes the model’s predicted probability of token  $y_t$  conditioned on the input  $x$  and the preceding tokens  $y_{<t}$ . We further normalize the perplexity score within each batch by dividing by the batch mean:

$$\widetilde{\text{PPL}}(x, y) = \frac{\text{PPL}(x, y)}{\frac{1}{B} \sum_{i=1}^B \text{PPL}(x^{(i)}, y^{(i)})}, \quad (9)$$

where  $B$  denotes the batch size, and  $\text{PPL}(x^{(i)}, y^{(i)})$  is the perplexity of the  $i$ -th sample in the batch.

For local uncertainty information weighting, we use the formulation from WPO (Zhou et al., 2024) as our local weighting strategy:

$$w(x, y) = \exp \left( \frac{1}{|y|} \sum_{t=1}^{|y|} \log \frac{\pi_{\theta}(y_t \mid x, y_{<t})}{\sum_{v \in \mathcal{V}} \pi_{\theta}(v \mid x, y_{<t})^2} \right), \quad (10)$$

For DPO, local weights are computed for both positive and negative pairs and multiplied to obtain the final weight, whereas for SFT, only the positive samples are used.

#### B.2. Hyperparameter Study

For SFT, the learning rate in Llama-3.2-1B-Instruct is chosen from  $\{2\text{e-}6, 5\text{e-}6, 7\text{e-}6, 1\text{e-}5\}$ , and in Llama-3.1-8B-Instruct from  $\{2\text{e-}5, 5\text{e-}5, 7\text{e-}5, 1\text{e-}4\}$ .

For DPO and KTO, we tune the temperature parameter  $\beta$  within the set  $\{0.01, 0.05, 0.1\}$ . In Llama-3.2-1B-Instruct, we search the learning rate in  $\{2\text{e-}7, 5\text{e-}7, 7\text{e-}7, 1\text{e-}6\}$ , while for Llama-3.1-8B-Instruct, the learning rate is selected from  $\{2\text{e-}6, 5\text{e-}6, 7\text{e-}6, 1\text{e-}5\}$ .

The baseline method and EAST share the same set of hyperparameters as the vanilla method to ensure a fair comparison. For EAST, we additionally search the exponent parameter  $a$  from the range  $\{-3, -2.5, -2, -1.5, -1.25, -1, -0.75, -0.5, 0.1, 0.2, 0.5, 0.7, 1, 1.5, 2, 2.5, 3\}$ .

For the ablation study, we report the average accuracy scores for  $a \in \{0.5, 1, 1.5\}$  across all three weighting methods.