

TRAINING LARGE LANGUAGE MODELS TO REASON IN PARALLEL WITH GLOBAL FORKING TOKENS

Sheng Jia^{1,2,3†}, Xiao Wang^{2*}, Shiva Prasad Kasiviswanathan^{2*}

¹University of Toronto, ²Amazon, ³Vector Institute

{sheng.jia@mail.utoronto.ca, waxao@amazon.com, kasivisw@gmail.com}

Code: <https://github.com/Sheng-J/SSFT>

Collection: shengjia-toronto/ssf

ABSTRACT

Although LLMs have demonstrated improved performance by scaling parallel test-time compute, doing so relies on generating reasoning paths that are both diverse and accurate. For challenging problems, the forking tokens that trigger diverse yet correct reasoning modes are typically deep in the sampling tree. Consequently, common strategies to encourage diversity, such as temperature scaling, encounter a worsened trade-off between diversity and accuracy. Motivated by this challenge, we treat parallel reasoning as a set-of-next-token-prediction problem, and incorporate a set-based global loss into Supervised Fine-Tuning (SFT) using self-supervised bipartite matching between our *global forking tokens* and unique reasoning traces. We observe that, while naive fine-tuning with multiple reasoning traces collapses these unique reasoning modes, our proposed method, Set Supervised Fine-Tuning (SSFT), preserves these modes and produces emergent global forking tokens. Experiments on multiple reasoning benchmarks show that our SSFT consistently outperforms SFT under both Pass@1 and Cons@ k metrics.

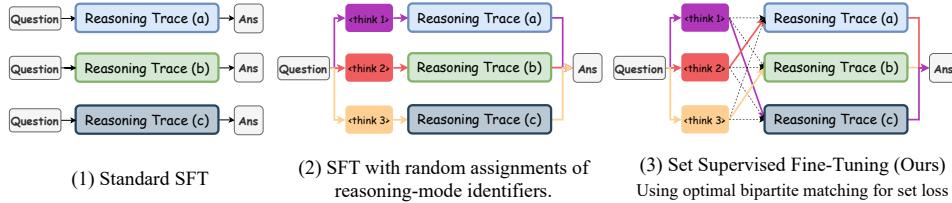


Figure 1: An illustration of different supervised fine-tuning methods that aim to instill parallel reasoning capabilities from diverse reasoning traces for the same question. Compared to (1) standard SFT and (2) SFT with randomly assigned parallel thinking identifiers, (3) Set-Supervised Fine-Tuning leverages a self-supervised bipartite matching process to learn to maximally differentiate the reasoning modes conditioned on distinct `<think i>`, for each question. Self-supervised matching prevents collapse of maximally distinct reasoning modes caused by ordering bias from randomly or manually assigning a reasoning mode identifier to a reasoning trace. The learned *global forking tokens* are considered emergent because we observe that similar reasoning modes over the traces for different questions cluster to match the same `<think, i>` at convergence, even though we do not manually identify such modes or incorporate any regularization term to maintain consistency in matching over different input prompts.

1 INTRODUCTION

Large language models have recently improved reasoning by allocating more test-time compute to generate more tokens before producing the final answer (OpenAI, 2025). However, extended

^{*}Equal advising.

[†]Work done during an internship at Amazon Science.

sequential scaling can lead to “overthinking”, where performance decreases beyond a certain sequence length (Ghosal et al., 2025; Chen et al., 2024a). To mitigate this, another scaling dimension based on repeated parallel sampling and aggregation (Wang et al., 2022; Brown et al., 2024) has shown success in further boosting reasoning performance. However, these methods rely on LLMs generating diverse yet correct solutions; as tasks become harder, a new mechanism is required to increase diversity. Recent work shows that only a minority of tokens in Chain-of-Thought reasoning (Wei et al., 2022) can act as forking tokens that lead to distinct reasoning modes (Wang et al., 2025b;a), so as the problem becomes harder and the generation becomes longer, it can become substantially harder to sample them. Also, common practices to encourage diversity, typically through temperature scaling, inherently entail a diversity-accuracy trade-off, as the forking tokens that trigger diverse yet correct reasoning modes are typically located deeply within the sampling tree. Moreover, recent theoretical work also shows that increasing the temperature alone does not necessarily guarantee greater diversity unless the model is explicitly trained for coverage. (Verine et al., 2025).

Building on these observations, we aim to leverage diverse reasoning traces to train for coverage (Guo et al., 2025; Google, 2025b). We introduce *global forking tokens* to initiate parallel reasoning traces and frame parallel reasoning as a *set prediction* problem. Specifically, given a question, an LLM, conditioned on a reserved set of tokens in a chosen ordering, generates M reasoning sequences in parallel, each aligned with one of M ground-truth reasoning traces. For each ordering, we compute the total autoregressive loss across the generated sequences. By enumerating all possible orderings, we identify the minimum loss, which defines the set language modeling loss, conditioned on the distinct forking tokens (Equation 5). This formulation naturally incorporates coverage into the training objective and is capable of learning *global forking tokens* that can serve as prompts to trigger reasoning modes that are both diverse and accurate. We operationalize this idea through our Set Supervised Fine-Tuning (SSFT) framework. Our main contributions are summarized below.

- We introduce global forking tokens and incorporate a set-based loss into SFT using bipartite matching between a set of reserved global forking tokens and diverse reasoning traces (Section 2; Figure 2). The goal is that, after training with SSFT, the model can initiate distinct reasoning modes when prompted with different tokens from this set, thus reducing the dependence on sampling forking tokens mid-generation (Section 2.3; Figure 3).
- We empirically demonstrate that, across multiple reasoning benchmarks, a model fine-tuned with SSFT outperforms a standard SFT model trained on the same distilled traces from diverse teacher models, improving Pass@1, Pass@ k , and Cons@ k (Section 3, Table 1 and Figure 4). We also find that naively fine-tuning using diverse reasoning traces without the set loss from bipartite matching can result in such control tokens initiating the same collapsed reasoning mode (Figure 6). In contrast, the global forking tokens learned by SSFT maximally differentiate reasoning modes (Figure 5).
- To facilitate training where different input prompts can use different numbers of parallel reasoning traces, which can cause difficulty in using gradient accumulation under distributed training, we also include a simple implementation based on the HuggingFace repository. Our implementation does not concatenate a variable number of reasoning traces along the sequence dimension and modify causal attention, which would require sequence parallelism and more complicated fused kernels to avoid out-of-memory for training with long reasoning traces (Appendix A.1).

2 LEARNING GLOBAL FORKING TOKENS VIA SET SUPERVISED FINETUNING

Background on Language Modeling with Reasoning. In language modeling, the goal is to train a model π_{θ} to approximate the joint distribution over a sequence of word tokens $\mathbf{x} = \{x_i\}_{i=1}^T \in \mathcal{V}^T$, where T is the sequence length, and each token is within a finite vocabulary set \mathcal{V} . An autoregressive model uses the chain rule to represent it as a product of conditionals on the preceding tokens:

$$\pi_{\theta}(\mathbf{x}) = \prod_{t=1}^T \pi_{\theta}(x_t | \mathbf{x}_{<t}) \quad (1)$$

This is known as *next-token-prediction* (NTP) (Radford et al., 2019). For reasoning tasks, we break the sequence of word tokens into: (1) an input prompt $\mathbf{x} = \{x_t\}_{t=1}^{T_x}$, (2) a special token (or a sequence

of tokens) g that initiates reasoning, and (3) a reasoning trace plus the final answer $\mathbf{r} = \{\mathbf{r}_t\}_{t=1}^{T_r}$. To simplify notation, we combine a verifiable answer and a reasoning trace. A reasoning model autoregressively generates a reasoning path and the final answer: $\pi_{\theta}(\mathbf{r}|\mathbf{x}, g) = \prod_{t=1}^{T_r} \pi_{\theta}(\mathbf{r}_t|\mathbf{x}, g, \mathbf{r}_{<t})$. To train a reasoning model, Supervised Fine-tuning (SFT) minimizes the negative log-likelihood of a ground-truth reasoning trace, i.e.,

$$\mathcal{L}(\theta) = -\mathbb{E}_{\mathbf{x}, \mathbf{r}} \left[\sum_{t=1}^{T_r} \log \pi_{\theta}(\mathbf{r}_t|\mathbf{x}, g, \mathbf{r}_{<t}) \right] \quad (2)$$

2.1 PARALLEL REASONING AS SET OF NEXT TOKEN PREDICTION

Problem Setup. In this paper, our goal is not only to instill new reasoning capabilities into a model, but also to ensure that prompting with a set of reserved special tokens, in parallel with a question, elicits distinct reasoning traces. We call these *global forking tokens* $\mathbf{g} := \{g^{(i)}\}_{i=1}^N$ instantiated as $\{\langle \text{think i} \rangle\}_{i=1}^N$ tags. We use $g^{(i)}$ interchangeably with $\langle \text{thinki} \rangle$, depending on context for clarity. We consider a setting with multiple sources of reasoning traces, obtained at low cost without human annotation by distilling from diverse teachers, sampling repeatedly, and potentially using rejection sampling to filter out low-quality or incorrect traces. We adopt this low-cost regime to highlight the effectiveness of our algorithm, though the method extends naturally to settings with well-annotated, human-labeled data. So our problem is to (1) learn to do a set of next-token predictions on multiple distinct yet correct reasoning traces $\mathbf{R} := \{\mathbf{r}^{(j)}\}_{j=1}^M$ in parallel for an input prompt \mathbf{x} , and (2) ensure that distinct global forking tokens can uniquely initiate these distinct traces.

To do this, we make a simple change to the NTP loss, which now has two requirements: **(1) Permutation-invariance:** It should not depend on the order of elements in \mathbf{R} and \mathbf{g} , so we don't penalize a trace that incurs high NTP loss under one global forking token if the model predicts it well when conditioned on another. **(2) No shared global forking token:** We want $\{g^{(i)}\}_{i=1}^N$ to uniquely initiate distinct reasoning traces, so this requirement prevents conditioning on the same $g^{(i)}$ when generating distinct traces given a question.

To satisfy these requirements, we formally incorporate a subproblem in language modeling: finding the minimum cost *bipartite matching configuration* between the left vertices $\{g^{(i)}\}_{i=1}^N$ and the right vertices $\{\mathbf{r}^{(j)}\}_{j=1}^M$ where the cost of each edge between a left vertex i and a right vertex j , is the NTP loss of $\mathbf{r}^{(j)}$ conditioned on $g^{(i)}$ and an input prompt \mathbf{x} . A *matching configuration* is a set of edges connecting the left and right vertices where no two edges share a common vertex. Without loss of generality, we assume $N \geq M$ to simplify our notation. The total cost involves all vertices on the smaller side of the bipartite graph, so this allows us to write the summation from 1 to $\min\{N, M\}$, which equals M under this assumption. We denote a matching configuration as a finite map $\sigma : \{1, \dots, M\} \rightarrow \{1, \dots, N\}$ such that $\sigma(j) = i \iff \mathbf{r}^{(j)}$ is paired with $g^{(i)}$. Let $\mathfrak{S}_P := \{\sigma_k\}_{k=1}^P$ denote all the $\binom{N}{M} \times M!$ configurations of a bipartite graph. The total cost of each configuration represents the compatibility between $\{\mathbf{r}^{(j)}\}_{j=1}^M$ and $\{g^{(i)}\}_{i=1}^N$ under this unique matching. Figure 2 shows an example of a matching configuration.

2.2 SSFT: MINIMIZING SET OF NTP LOSSES UNDER OPTIMAL BIPARTITE MATCHING

Under this formulation, we propose Set Supervised Fine-Tuning (SSFT), which performs two operations at each training step: (1) find the minimum-cost bipartite matching that is optimal for uniquely initiating different reasoning traces (2) and then minimize the NTP losses under the matching configuration to instill diverse reasoning modes conditioned on the matched global forking tokens. We show our implementation in Algorithm 1.

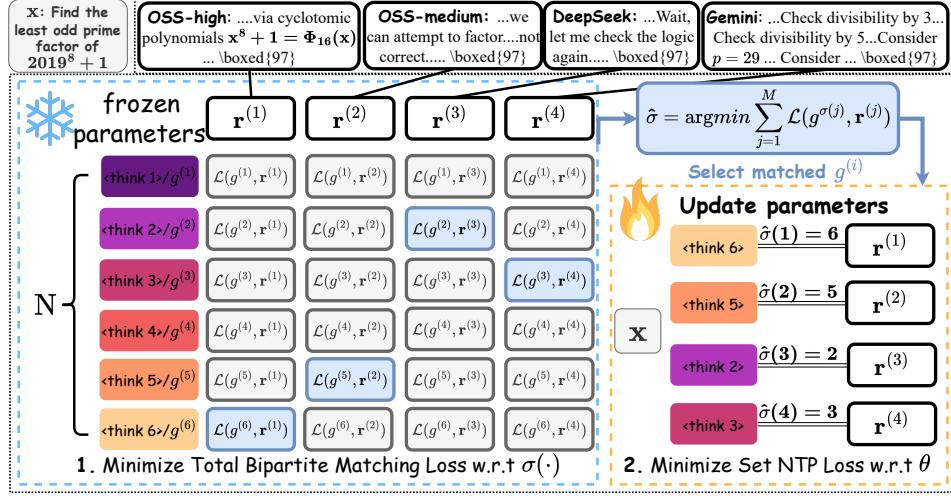


Figure 2: An illustration of one SSFT training step. **Step 1:** We first construct the cost matrix by evaluating all pairwise combinations: for each $\mathbf{r}^{(j)} \in \{\mathbf{r}^{(1)}, \mathbf{r}^{(2)}, \mathbf{r}^{(3)}, \mathbf{r}^{(4)}\}$ and each $\mathbf{g}^{(i)} \in \{\mathbf{g}^{(1)}, \mathbf{g}^{(2)}, \mathbf{g}^{(3)}, \mathbf{g}^{(4)}, \mathbf{g}^{(5)}, \mathbf{g}^{(6)}\}$, we compute the NTP loss of $\mathbf{r}^{(j)}$ conditioned on $\mathbf{g}^{(i)}$ (Equation (4)). Then we use Hungarian algorithm to find $\hat{\sigma}$ that minimizes the total bipartite matching cost. Here, this minimum is the sum of the losses highlighted in blue, which means $\hat{\sigma} = \{(\mathbf{g}^{(6)}, \mathbf{r}^{(1)}), (\mathbf{g}^{(5)}, \mathbf{r}^{(2)}), (\mathbf{g}^{(2)}, \mathbf{r}^{(3)}), (\mathbf{g}^{(3)}, \mathbf{r}^{(4)})\}$. **Step 2:** We optimize θ by backpropagating the set of NTP losses for $\mathbf{r}^{(j)}$, each conditioned on $\mathbf{g}^{(\hat{\sigma}(j))}$. This is the Hungarian loss in Equation 5.

For the first step, we compute all the entries in the cost matrix such as the one in Figure 2 and then apply the Hungarian algorithm (Kuhn, 1955) to efficiently find the optimal $\hat{\sigma}$:

$$\hat{\sigma} = \operatorname{arg} \min_{\sigma \in \mathfrak{S}_P} \sum_{j=1}^M \mathcal{L}_{\text{matching}}\left(\mathbf{g}^{(\sigma(j))}, \mathbf{r}^{(j)}\right), \text{ where} \quad (3)$$

$$\mathcal{L}_{\text{matching}}\left(\mathbf{g}^{(i)}, \mathbf{r}^{(j)}\right) = -\operatorname{sg}\left(\frac{1}{T_r} \sum_{t=1}^{T_r} \log \pi_{\theta}\left(\mathbf{r}_t^{(j)} | \mathbf{x}, \mathbf{g}^{(i)}, \mathbf{r}_{<t}^{(j)}\right)\right) \quad (4)$$

As noted in Equation 4, each matching cost in Equation 3 is the negative log-likelihood of $\mathbf{r}^{(j)}$ conditioned on $\mathbf{g}^{(\sigma(j))}$ under the current model parameters. Here, $\operatorname{sg}(\cdot)$ is stop-gradient, as the matching process is done by discrete optimization w.r.t. σ , so we can save VRAM by not storing intermediate activations. We explicitly indicate that length normalization is done to remove biases toward trace length, so that the matching is driven by semantic content.

After solving $\hat{\sigma}$ for each (\mathbf{x}, \mathbf{R}) , our second step optimizes model parameters θ by backpropagating the Hungarian loss in Equation 5. In Algorithm 1, the expectation is replaced by its sample mean over pairs of $(\mathbf{x}, \{\mathbf{r}^{(j)}\}_{j=1}^M)$ in a mini-batch. In practice, we may use only the first $L < T_r$ tokens in Equation 4 to compute the matching cost and find $\hat{\sigma}$ when the training dataset is such that these early tokens reveal sufficient differences. However, we always optimize the Hungarian loss for the full T_r length.

$$\mathcal{L}_{\text{Hungarian}}(\theta) = -\mathbb{E}_{\mathbf{x}, \mathbf{R} \sim \mathcal{D}} \left[\sum_{j=1}^M \sum_{t=1}^{T_r} \log \pi_{\theta}\left(\mathbf{r}_t^{(j)} | \mathbf{x}, \mathbf{g}^{(\hat{\sigma}(j))}, \mathbf{r}_{<t}^{(j)}\right) \right] \quad (5)$$

Remarks. The resulting model is not the same as a simple routing of the models independently trained with the nonoverlapping subsets of these traces. Firstly, SSFT allows positive transfer in representation learning within $\{\mathbf{r}^{(j)}\}_{j=1}^M$ even though they are matched to different $\{\mathbf{g}^{(i)}\}_{i=1}^N$. Secondly, it is not optimal to distill reasoning traces from the same fixed sources for every question if the goal is to maximize both diversity and correctness. Some traces can be filtered by rejection sampling. Our algorithm supports a variable number of target reasoning traces across training steps,

with sources that may also change for each x . Thirdly, even if the two sets of traces, $\{\mathbf{r}_a^{(j)}\}_{j=1}^M$ for \mathbf{x}_a and $\{\mathbf{r}_b^{(j)}\}_{j=1}^M$ for \mathbf{x}_b , are from the same sources, their optimal configurations $\hat{\sigma}_a$ and $\hat{\sigma}_b$ can still vary because a teacher model can reason differently under different questions. Lastly, we reserve more global forking tokens than the maximum number of traces ($N > M$), and empirically observe that all the forking tokens are being matched throughout the process. This is because the extra forking tokens can maximally intra-differentiate similar traces.

2.3 INFERENCE WITH LEARNED GLOBAL FORKING TOKENS

We discuss the inference protocols with N global forking tokens.

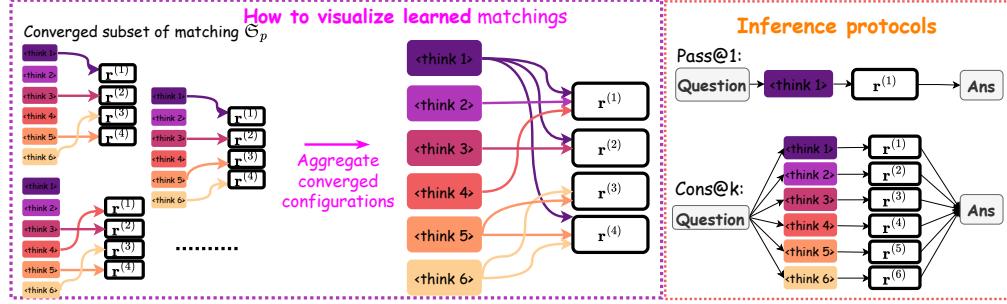


Figure 3: **(Left)** \mathfrak{S}_p is the subset of bipartite matching configurations that are still computed optimal towards the end of training. These can be individually visualized as a bipartite **(Middle)** Learned matchings by SSFT-32B in Exp 3, obtained by aggregating all edges in \mathfrak{S}_p . **(Right)** At test time, for Pass@1, we prompt with $g^{(i^*)}$ that has the most connected edges. For Cons@ k , we prompt i -th parallel generation with $\langle \text{think } (i \% N) \rangle$.

Inference Process (Cons@ k). Our inference protocol with parallel test-time compute is to prompt i -th response with $\langle \text{think } (i \% N) \rangle$ and then do majority voting on their answers. KV prefix caching accelerates the generations.

Inference Process (Pass@1). When aggregation is not allowed, we use $g^{(i)}$ that reasons with more flexibility. Inspired by enumerating dissimilar bipartite matchings to reveal node-level variation (Blumenthal et al., 2022), we choose the learned $g^{(i)}$ with largest coverage. Note that this token emerges automatically thanks to SSFT, and the other emerging $\{g^{(i)}\}_{i=1}^N$ still contribute to representation learning and improving Cons@ k with distinct reasoning modes. Concretely, each time an optimal matching $\hat{\sigma} \in \{\sigma_k\}_{k=1}^P$ is computed, we increment a count $c(\sigma_k)$; empirically, only a finite subset $\mathfrak{S}_p := \{\sigma_k\}_{k=1}^p \subseteq \mathfrak{S}_P$ continues to accumulate mass late in training, indicating the stable learned matchings. We then take the union of their edges, and select $g^{(i^*)}$ that matched to the largest number of distinct traces based on Equation 6 for Pass@1. Figure 3 shows an example of the matchings learned by aggregating all edges in \mathfrak{S}_p . More details about \mathfrak{S}_p ¹ are in Appendix A.2.

$$i^* = \arg \max_i \left| \bigcup_{\sigma \in \mathfrak{S}_p} \{j | \sigma(j) = i\} \right| \quad (6)$$

3 EXPERIMENTS

We address the following research questions through experiments: **(RQ1):** In terms of Pass@1 and Cons@ k accuracy, how does a model trained with SSFT perform on reasoning benchmarks? **(RQ2):** Does finding the optimal bipartite matching matter in reasoning performance? **(RQ3):** Does training with diverse reasoning traces yield better accuracy and coverage under SSFT compared to standard SFT with temperature scaling? **(RQ4):** Does prompting with distinct $\{g^{(i)}\}_{i=1}^N$ genuinely make a model generate diverse reasoning traces? **(RQ5):** Is the performance gain from SSFT conditioned on the traces generated by our procedure and on a high-quality small dataset?

¹We choose the subscripts p and P to emphasize that \mathfrak{S}_p is a subset of \mathfrak{S}_P .

3.1 EXPERIMENT SETUP

Training Dataset. We use the 1,000 questions from s1k dataset [Muennighoff et al. \(2025\)](#). In addition to the R1 ([Guo et al., 2025](#)) and Gemini Flash ([Google, 2024](#)) traces provided by s1, we also use Claude Opus 4.0/4.1 ([Anthropic, 2025](#)) and GPT-OSS-120B ([Agarwal et al., 2025](#)) with high and medium reasoning effort to obtain a pool of distilled targets for the 1,000 questions. For each question, we generate two traces per source to populate the pool. We then sample four traces from this pool. We call this s1k-4mixed-reasoning dataset.

Training Details. We fine-tune Qwen2.5-32B-Instruct ([Yang et al., 2025a](#)) for six epochs with a context length of 32,768. We reserve $N = 6$ global forking tokens and use $M = 4$ targets per question. To find the optimal bipartite matching for each input prompt, we consider only the first 1,000 tokens when computing the matching cost in Equation 4 for computational efficiency. We call this model SSFT-32B. We also include SSFT but choose a random bipartite matching at each step to fine-tune SSFT-32B (random σ). Exact details on the pool of diverse distillation targets and selection procedure, as well as training hyperparameters, are provided in Appendix A.3.

Baselines. All of our baselines use Qwen2.5-32B-Instruct as their base model, and only train on the 1k questions. Our baselines include two groups: **(Single-Target Δ)** models trained with one trace per question and **(Multi-Target \star)** models trained with four traces per question. For **(Single-Target)**, we include s1.1-32B ([Muennighoff et al., 2025](#)), which uses 1k DeepSeek-R1 traces. We also fine-tuned an SFT-OSS-distill-32B baseline that trains only on the 1k GPT-OSS traces with high reasoning effort, as these traces achieved the highest correctness on the 1k questions based on an evaluation by Claude 3.5 Sonnet comparing each attempt against the reference answer. For **(Multi-Target)**, we use our s1k-4mixed-reasoning to fine-tune SFT-mixed-distill-32B using standard SFT with one `<think>` token, duplicating each question and treating the four traces as four individual data points. We also include Multiverse-32B ([Yang et al., 2025b](#)), which prompts Gemini 2.5-Pro ([Google, 2025a](#)) to transform 1k sequential CoTs into parallel CoTs as their training data.

Evaluation Setup. Our evaluation tasks consist of AIME24/AIME25 ([Ye et al., 2025](#)), MATH-500 ([Hendrycks et al., 2021](#)), and GPQA-Diamond ([Rein et al., 2024](#)). We use LightEval ([Habib et al., 2023](#)) as our evaluation framework with generation configurations: temperature=0.7 used in ([Guha et al., 2025](#)), top_p=0.95, max length=32768. For Pass@1 accuracy without any parallel test-time compute, we select learned $g^{(1)}$ for SSFT-32B and $g^{(4)}$ for SSFT-32B (random σ) based on Equation 6. For each Pass@1 accuracy, we compute the average performance over 32 generations. For Cons@6, which applies each of the six global forking tokens once in our method and uses six generations for the baselines, we compute the average over 11 sets of generations to reduce variance in the results. We refer to this as *Average of Native Cons@6*, which corresponds to the Pass@1 metric with parallel generations in concurrent work ([Wen et al., 2025](#)). Appendix A.4 presents an example of the parallel generations by SSFT-32B.

3.2 EVALUATING SSFT ON REASONING BENCHMARKS

For **RQ1**, we see in Table 1 that SSFT delivers the best Pass@1 accuracy, 64.06 on AIME24 and 58.13 on AIME25, outperforming SFT-mixed-32B by 8.33% and 6.57%, respectively. We also observe consistent improvements on all four tasks under parallel test-time compute at two scales, Cons@6 and Cons@32, over SFT-mixed-32B, which was trained on the same reasoning traces. Some notable results are Cons@6 = 73.94%, Cons@32 = 86.67% on AIME25. To answer **RQ2**, we observe consistent improvements over SSFT-32B (random σ), with especially strong gains at Cons@6 on AIME25, where effectiveness with few parallel generations is critical. As shown later in Figures 5 and 6, optimal bipartite matching is essential for preventing collapsing reasoning modes.

For **RQ3**, we compare our method against SFT-mixed-32B under various k in Pass@ k accuracy with 32 generations to assess generation coverage. Figure 4 shows that SSFT achieves higher coverage across nearly all values of k . SFT-mixed-32B requires more allowed attempts and higher temperature to match the coverage of SSFT at the cost of lowering its Pass@1 and Cons@6 accuracy.

	AIME 2024	AIME 2025	MATH-500	GPQA-D	Average
<i>Pass@1: Average performance of individual generations</i>					
Qwen2.5-32B-Instruct \triangle	15.80	10.40	80.40	47.00	38.40
s1.1-32B \triangle	54.79	44.27	92.16	62.12	63.34
Multiverse-32B \star	53.80	45.80	91.80	60.70	63.03
SFT-OSS-distill-32B \triangle	57.82	48.75	89.54	60.06	64.04
SFT-mixed-distill-32B \star	55.73	51.56	88.36	57.50	63.29
SSFT-32B (random σ) \star	61.77	55.10	89.95	62.28	67.28
SSFT-32B \star	64.06	58.13	90.02	60.39	68.15
<i>Average of Native Cons@6: Average performance of majority voting with 6 parallel generations</i>					
s1.1-32B \triangle	70.30	53.33	95.60	61.45	70.17
SFT-OSS-distill-32B \triangle	72.12	65.45	95.47	61.52	73.64
SFT-mixed-distill-32B \star	72.42	70.91	92.10	57.32	73.19
SSFT-32B (random σ) \star	73.03	67.58	95.67	61.87	74.54
SSFT-32B \star	75.45	73.94	96.47	63.05	77.23
<i>Cons@32: Majority voting performance with large number of parallel generations</i>					
s1.1-32B \triangle	73.33	63.33	94.80	60.61	73.02
SFT-OSS-distill-32B \triangle	76.66	73.33	96.00	61.60	76.90
SFT-mixed-distill-32B \star	80.00	73.33	96.20	60.61	77.54
SSFT-32B (random σ) \star	80.00	80.00	95.60	62.63	79.56
SSFT-32B \star	83.33	86.67	96.80	61.62	82.11

\triangle indicates training with single-target data and \star indicates training with multi-target data.

Table 1: Performance of SSFT compared to baselines on four reasoning tasks, reported at Pass@1, Cons@6, and Cons@32. SSFT selects `<think1>` for Pass@1 and replaces 6 generations with 6 generations prompted by distinct `<thinki>` for Cons@ k . We observe consistent improvements over (i) SFT-OSS-distill-32B, which uses the 1k OSS-high traces; (ii) SFT-mixed-distill-32B, which uses the four mixed traces but treats them as individual data; and (iii) SSFT-32B (random σ), which trains using Equation 5 but with a randomly chosen σ .

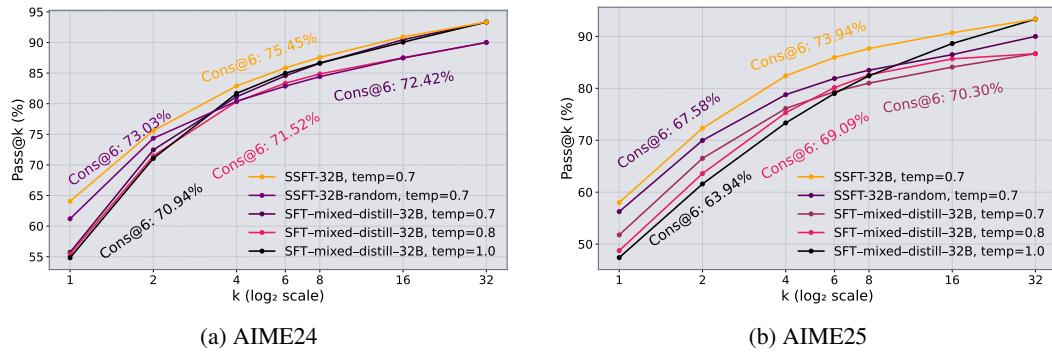


Figure 4: Coverage of SSFT compared to SFT-mixed-distill-32B with temperature scaling, reported at Pass@ k . For convenience, we also report the Cons@6 accuracy next to each line. In AIME25, SFT-mixed-distill-32B needs to raise the inference temperature to 1 and use more attempts to match the coverage at the cost of lowering its Pass@1 and Cons@6 accuracy, further widening the gaps.

3.3 EVALUATING PARALLEL REASONING DIVERSITY AND LEARNED MATCHINGS

Addressing **RQ4**, we show that our global forking tokens genuinely initiate distinct reasoning traces and offer a new mechanism for leveraging test-time compute. We call these tokens emergent because we do not manually identify distinct reasoning modes over the traces for different questions. For example, `<think1>` could lead to longer reasoning than `<think5>` for one question, but this may not be consistent for other similar questions. We do not add any auxiliary objective to skew the original language modeling objective.

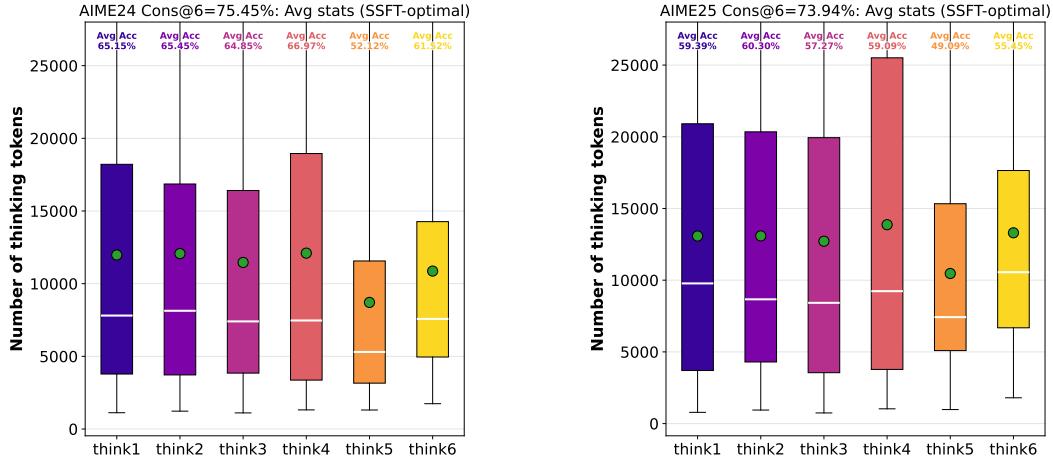


Figure 5: (SSFT, optimal matching). Distribution of thinking-token counts and average performance on AIME24 (left) and AIME25 (right) prompted by a distinct $\langle \text{think1}, \dots, \text{think6} \rangle$. Each accuracy is averaged over 11 generations, shown above each whisker. We observe clear specialization of reasoning modes. $\langle \text{think1}, \dots, \text{think4} \rangle$ trigger longer traces, better suited to many AIME problems, and these surpass all the average accuracy under SFT without bipartite matching in Figure 6. $\langle \text{think5} \rangle$ and $\langle \text{think6} \rangle$ favor concise reasoning, which can be suboptimal for some AIME tasks, yet still raise Cons@6 beyond all baselines by adding diversity. The consistency of these length distributions across AIME24 and AIME25 indicates the difference is not due to randomness, and that these think tags truly initiate distinct yet consistent reasoning modes.

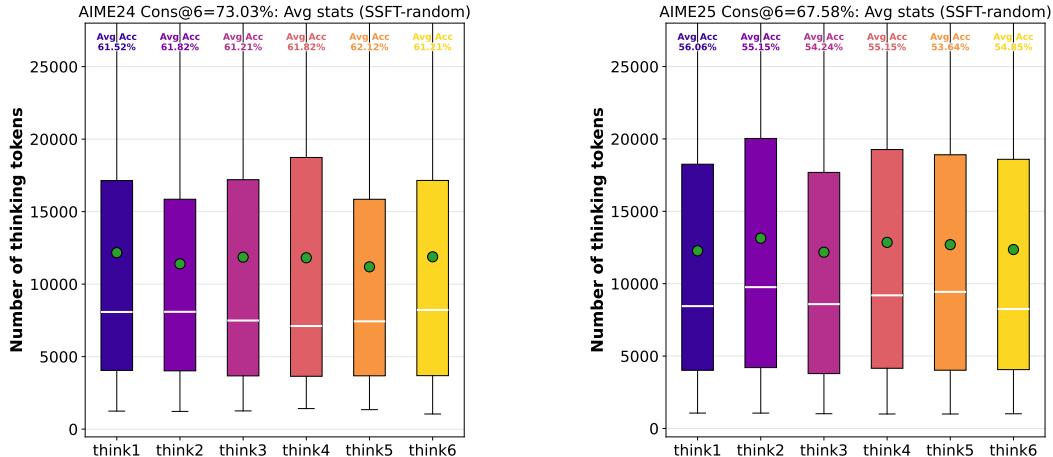


Figure 6: (SSFT, random matching). Distribution of thinking-token counts and average performance on AIME24 (left) and AIME25 (right) prompted by a distinct $\langle \text{think1}, \dots, \text{think6} \rangle$. Each accuracy is averaged over 11 generations, shown above each whisker. Without optimal bipartite matching during SFT, the think tags initiate reasoning traces with similar length distributions and achieve similar average performance. Moreover, the slight length differences are not consistent across AIME24 and AIME25, suggesting randomness rather than genuine specialization; the tags therefore collapse to similar modes, failing to preserve the diverse reasoning modes in the training data. Both the average performance and Cons@6 fall behind those from any think tag that initiates long reasoning under SSFT in Figure 5.

Emerging Diverse Reasoning Modes. Using the Cons@6 results in Table 1, we form six sets of generations, each prompted by a distinct $g^{(i)}$. For each set, we show the average accuracy and the distribution of thinking-token counts: Figure 5 for SSFT with optimal bipartite matching and Figure 6 for SSFT with random matching. **Reasoning length:** Length partially indicates the diversity in reasoning, and we see clear differences for SSFT with optimal matching, despite the absence of hand-crafted matching rule or information about these traces. The consistency of these distributions

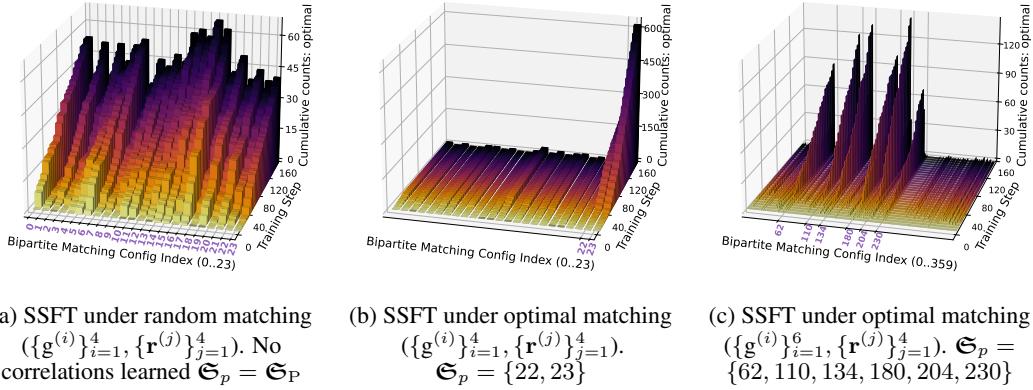


Figure 7: Cumulative counts of $\sigma_k \in S_p$ computed as optimal over training. Note that S_P and S_p are defined in Sections 2.1 and 2.3, respectively. Front axis: matching configuration index k . Depth: training step t . Bar height is the cumulative counts. These are the evolution of matchings during training 3 Qwen-32B-Instruct models under 3 bipartite matching settings. In this case study, the $\{r^{(j)}\}_{j=1}^4$ are always (GPT-OSS-high, GPT-OSS-med, R1, Gemini) for each question. Note that *Random matching* method does not minimize Eqn 5 under optimal matching, but we track it. We observe $S_p = S_P$ with random matching, meaning no correlations learned. But by optimizing Hungarian loss, we see the emergence of $S_p \subset S_P$.

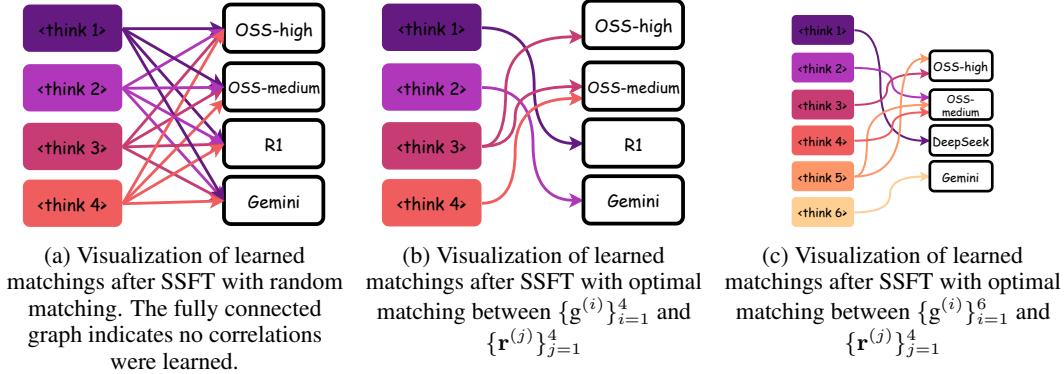


Figure 8: We visualize the learned matchings by connecting edges in the subset of configurations S_p that still accumulate mass late in training (Figure 7). All the three models in (a), (b), and (c) are fine-tuned using GPT-OSS-high, GPT-OSS-medium, R1, and Gemini to distill reasoning traces for every question, enabling clearer interpretation.

across AIME24 and AIME25 indicates the differences is not from randomness, whereas randomly assigning a $\langle \text{think}_i \rangle$, as in concurrent work (Wen et al., 2025), does not yield clear or consistent differences in reasoning length, as shown in Figure 6. **Performance:** After finetuning with random matching, prompting with a distinct $\langle \text{think}_i \rangle$ shows no meaningful impact ($\approx 61\%$ on AIME24 and $\approx 55\%$ on AIME25). With optimal matching, SSFT elicits distinct reasoning modes initiated by $\langle \text{think}_1 \rangle, \dots, \langle \text{think}_4 \rangle$ that reach around 65% on AIME24 and $\geq 59\%$ on AIME25, with different lengths. Although $\langle \text{think}_5 \rangle$ and $\langle \text{think}_6 \rangle$ are weaker due to shorter reasoning modes, the average of these and, especially, Cons@6 performance improve consistently with them.

Visualization of Learned Matchings. We observe in Figure 3 that some $g^{(i)} \in g$ from SSFT-32B have unique configuration of matched edges with $\{r^{(j)}\}_{j=1}^4$. This is a positive indication that $\{g^{(i)}\}_{i=1}^6$ are likely to initiate distinct reasoning modes. We hypothesize that $\{g^{(i)}\}_{i=1}^N$ can still yield a unique edge-matching configuration even if a subset of $\{r^{(j)}\}_{j=1}^4$ are difficult to distinguish.

For interpretability, we fine-tune using the same four teacher models for each x : GPT-OSS high, medium, R1, and Gemini. We call this dataset s1k-4teachers-reasoning dataset, and ask whether

SSFT associates a unique $g^{(i)}$ to the Gemini and R1 traces which have easily identifiable reasoning patterns, and then have the rest matched to OSS-high/OSS-medium traces in different ways (i.e. only matched to OSS-high, only matched to OSS-med, and matched to both). We track source indices only for evaluation; the model still receives an unordered set of traces with no source information for each question. We study three bipartite matching settings with SSFT. Figure 7 shows the evolution of matchings learned under these three hyperparameters: (a) random matching with four reserved $g^{(i)}$, (b) optimal matching with four reserved $g^{(i)}$, and (c) optimal matching with six reserved $g^{(i)}$. Initially, all the configurations $\sigma \in \mathfrak{S}_P$ accumulate mass as there is no correlations between g and $\{r^{(j)}\}_{j=1}^M$. Figure 7a shows SSFT under random matching does not shrink the size of configurations computed as optimal, meaning that no correlations are learned between $\{g^{(i)}\}_{i=1}^4$ and $\{r^{(j)}\}_{j=1}^4$. By contrast, Figure 7b and Figure 7c show the emergence of only a strict subset of matching configurations in \mathfrak{S}_P . This indicates some correlations between $\{g^{(i)}\}_{i=1}^N$ and $\{r^{(j)}\}_{j=1}^M$ are indeed learned through SSFT. We first visualize the learned matchings of the model with four $g^{(i)}$ in Figure 8b. We observe that $g^{(1)}$ and $g^{(2)}$ are uniquely matched to the R1 and Gemini traces, showing that SSFT can indeed uniquely associate $g^{(i)}$ to sufficiently diverse reasoning traces. Now to confirm our previous hypothesis, we see the unique learned matchings $(g^{(3)}, (\text{OSS-high, OSS-med}))$, $(g^{(4)}, \text{OSS-med})$. Furthermore, by connecting all the edges in \mathfrak{S}_P from SSFT with 6 forking tokens (Figure 7c), we also see unique learned matchings $(g^{(3)}, \text{OSS-high})$, $(g^{(4)}, \text{OSS-med})$, $(g^{(5)}, (\text{OSS-med, OSS-high}))$ in Figure 8c. This confirms that the global forking tokens can identify unique correlations even among highly similar traces.

3.4 ABLATION STUDY: REMOVING HIGH QUALITY SMALL DATASET

To test whether our empirical gains are conditioned on the traces generated by our procedure and on a highly optimized small dataset (Muennighoff et al., 2025), we fine-tune on a public dataset that already provides sufficient reasoning traces per question: the 93k math set of Face (2025) (2-4 traces per question). Because this dataset has been successful for fine-tuning Qwen2.5-Math-7B, we adopt that base model and compare SSFT against SFT trained on all available traces. Details on the hyperparameters are in Appendix A.3.4. Addressing RQ5, Table 2 shows consistent improvements in both Pass@1 and Cons@32. The results indicate the SSFT is effective for larger sized public dataset with less diverse reasoning traces.

Model	AIME 2024		AIME 2025		MATH-500		GPQA-D	
	Pass@1	Cons@32	Pass@1	Cons@32	Pass@1	Cons@32	Pass@1	Cons@32
Qwen2.5-Math-7B-Instruct	10.42	20.00	9.48	23.33	81.87	87.40	30.29	30.30
SFT-OpenR1-93k-7B	46.15	66.67	34.17	50.00	86.62	90.20	46.35	47.98
SSFT-OpenR1-93k-7B	51.25	73.33	35.52	56.66	89.74	93.60	46.86	48.90

Table 2: Performance of SSFT versus SFT trained solely on publicly available distillation targets. The setup uses the 93k math questions from Face (2025) with Qwen2.5-Math-7B as the base model. SFT-OpenR1-93k-7B uses the same distillation targets as SFT-mixed-distill-32B in Table 1.

4 RELATED WORK

Test-time Scaling. There has been a surge of work fine-tuning LLMs to reason longer, using reinforcement learning for frontier models (OpenAI, 2024; Shao et al., 2024; xAI, 2025; Yang et al., 2025a) and supervised fine-tuning for smaller ones (Muennighoff et al., 2025; Hu et al., 2025). These methods enable LLMs to improve reasoning by allocating more test-time compute to sequential, iterative refinement such as self-reflection (Guo et al., 2025; Liu et al., 2025). However, extended sequential reasoning can be more sensitive to the order of reasoning steps and may result in failures (Chen et al., 2024b), and performance can start to degrade beyond a certain length due to “overthinking” (Ghosal et al., 2025). Our goal is to study the effective use of diverse reasoning traces to fine-tune small language models, essential for agentic AI (Belcak et al., 2025). Related to our work on aggregating diverse reasoning modes, Yu et al. (2025); Zheng et al. (2025a) proposed incorporating three specific reasoning modes: natural language, code-based, and symbolic reasoning.

Parallel Reasoning. Parallel scaling methods such as self-consistency (Wang et al., 2022) and Best-of-N (Lightman et al., 2023) improve LLM performance by generating multiple reasoning paths in parallel and aggregating them. These methods fundamentally require choosing a temperature that can generate diverse reasoning paths, but a recent theoretical work shows that increasing temperature can sometimes fail to increase diversity if language models are not trained towards coverage (Verine et al., 2025). Other search-based methods such as Monte Carlo tree search (MCTS) (Zhang et al., 2024) and Tree of Thoughts (ToT) (Yao et al., 2023) apply heuristic-guided search with an external verifier to do more deliberate search to increase the coverage. However, their dependence on heuristics and domain-specific knowledge can limit their applicable tasks. Regarding training LLMs with parallel reasoning traces, Yang et al. (2025b) proposes training with parallel CoTs decomposed from sequential CoTs, and our concurrent work Wen et al. (2025) proposes to train with multiple reasoning traces distilled from teacher models. These works show native parallel scaling can surpass sequential scaling within certain token limits. Although we also show improvements in reasoning performance by simply fine-tuning with parallel reasoning traces, we empirically highlighted the limitation of manually assigning parallel think tags to target reasoning traces. If human assignment or random enforcement routes mismatched traces to a tag with a unique reasoning mode learned, the tag’s semantics blur due to bias in the ordering of parallel reasoning traces. This prevents tags from initiating distinct reasoning modes. We are the only work that leverages parallel think tags to construct a set-prediction loss via bipartite-matching optimization between tags and reasoning traces, turning the tags into global forking tokens. Empirically, we also improve Pass@1 and Cons@ k over baselines fine-tuned with parallel reasoning traces on the same dataset, whether on subsets or the full set. Another recent work (Zheng et al., 2025b) encourages the use of parallel thinking tags by slightly upweighting rewards for correct trajectories that employ such tags during reinforcement learning. However, the cold-start SFT that learns parallel thinking from diverse reasoning traces does not incorporate a set-prediction loss, which is necessary to prevent collapse in reasoning diversity.

Set-based Global Loss in Deep Learning. DETR introduces end-to-end object detection with a set global loss (Carion et al., 2020; Minderer et al., 2022), whose success in parallel bounding-box prediction inspires our approach. We are the first to extend this to language modeling: While DETR predicts a set of single tokens in parallel to match a list of bounding boxes, we predict a set of sequential thinking tokens, initiated by global forking tokens, and evaluate the matchings using autoregressive losses. We adapt set-based loss with autoregressive (AR) models, rather than diffusion-based models (Zhao et al., 2025; Nie et al., 2025; Ho et al., 2020), because AR models achieve superior reasoning performance. Several innovations are needed to make this loss practical for parallel reasoning with AR models. For example, to populate the cost matrix we need $N \times M$ forward passes, one for each combination of N think tags and M reasoning traces, whereas DETR requires only N forward passes for its N queries, since the same bounding-box predictions can be reused to compute matching costs against different ground-truth boxes. We therefore emphasize efficient training implementation in our extension, increasing training time by only 6.6% compared to standard SFT (Appendix A.1).

5 CONCLUSION

In this work, we demonstrate that diverse reasoning traces can be leveraged to learn global forking tokens which serve as prompts to initiate distinct reasoning modes that are both diverse and accurate. We propose Set Supervised Fine-Tuning (SSFT), which employs bipartite matching between reserved global forking tokens and diverse reasoning traces to compute a set-based language modeling loss. We show that models trained with SSFT yield improvements in both Pass@ k and Cons@ k accuracy over various k compared to standard SFT with temperature scaling. The proposed training method also improves the usage of multi-target data and avoids collapsing distinct reasoning modes, which is especially helpful in the context of distilling from multiple teacher models. In future work, we plan to further investigate how performance scales with a larger number of distillation targets and reserved forking tokens, which jointly determine the size of the underlying bipartite graph. In addition, we will apply RFT to a model pre-trained with SSFT to incentivize optimal use of the learned global forking tokens across a broader suite of tasks without requiring additional high-quality reasoning traces.

ACKNOWLEDGMENT

This research was supported by compute resources provided by Amazon Web Services.

REFERENCES

Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025.

Anthropic. Claude 4.1 Opus. <https://claude.ai>, 2025. Generative AI Chatbot. Response to a prompt on [Date of Access].

Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. Small language models are the future of agentic ai. *arXiv preprint arXiv:2506.02153*, 2025.

David B. Blumenthal, Sébastien Bougleux, Anton Dignös, and Johann Gamper. Enumerating dissimilar minimum cost perfect and error-correcting bipartite matchings for robust data matching. *Inf. Sci.*, 596:202–221, 2022. URL <https://doi.org/10.1016/j.ins.2022.03.017>.

Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*, 2024.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pp. 213–229. Springer, 2020.

Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qizhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024a.

Xinyun Chen, Ryan Andrew Chi, Xuezhi Wang, and Denny Zhou. Premise order matters in reasoning with large language models. In *Forty-first International Conference on Machine Learning*, 2024b. URL <https://openreview.net/forum?id=4zAHgkiCQg>.

Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023.

Hugging Face. Open r1: A fully open reproduction of deepseek-r1, January 2025. URL <https://github.com/huggingface/open-r1>.

Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Yifu Lu, Mengdi Wang, Dinesh Manocha, Furong Huang, Mohammad Ghavamzadeh, and Amrit Singh Bedi. Does thinking more always help? understanding test-time scaling in reasoning models. *arXiv preprint arXiv:2506.04210*, 2025.

Google. Gemini 2.0 flash thinking mode (gemini-2.0-flash-thinking-exp-1219), December 2024. URL <https://cloud.google.com/vertex-ai/generative-ai/docs/thinking-mode>.

Google. Gemini 2.5 (gemini-2.5-pro-preview). <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>, 2025a. Accessed: 2025-04-22.

Google. Gemini 2.5: Our most intelligent ai model. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/#gemini-2-5-thinking>, March 2025b.

Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, et al. Openthoughts: Data recipes for reasoning models. *arXiv preprint arXiv:2506.04178*, 2025.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Nathan Habib, Clémentine Fourrier, Hynek Kydlíček, Thomas Wolf, and Lewis Tunstall. Lighteval: A lightweight framework for llm evaluation, 2023. URL <https://github.com/huggingface/lighteval>.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Pin-Lun Hsu, Yun Dai, Vignesh Kothapalli, Qingquan Song, Shao Tang, Siyu Zhu, Steven Shimizu, Shivam Sahni, Haowen Ning, and Yanning Chen. Liger kernel: Efficient triton kernels for llm training. *arXiv preprint arXiv:2410.10989*, 2024.

Xiao Hu, Xingyu Lu, Liyuan Mao, YiFan Zhang, Tianke Zhang, Bin Wen, Fan Yang, Tingting Gao, and Guorui Zhou. Why distillation can outperform zero-rl: The role of flexible reasoning. *arXiv preprint arXiv:2505.21067*, 2025.

H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. doi: <https://doi.org/10.1002/nav.3800020109>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109>.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step, 2023. URL <https://arxiv.org/abs/2305.20050>.

Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European conference on computer vision*, pp. 728–755. Springer, 2022.

Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*, 2025.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025.

OpenAI. Learning to reason with llms, September 2024. URL <https://openai.com/index/learning-to-reason-with-llms/>.

OpenAI. Introducing openai o3 and o4-mini. <https://openai.com/index/introducing-03-and-04-mini/>, 2025.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. In *OpenAI Blog*, 2019. URL <https://api.semanticscholar.org/CorpusID:160025533>.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

Alexandre Verine, Florian Le Bronnec, Kunhao Zheng, Alexandre Allauzen, Yann Chevaleyre, and benjamin negrevergne. Improving diversity in language models: When temperature fails, change the loss. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=RsyMfsqzeG>.

Haozhe Wang, Qixin Xu, Che Liu, Junhong Wu, Fangzhen Lin, and Wenhui Chen. Emergent hierarchical reasoning in llms through reinforcement learning. *arXiv preprint arXiv:2509.03646*, 2025a.

Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025b.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Hao Wen, Yifan Su, Feifei Zhang, Yunxin Liu, Yunhao Liu, Ya-Qin Zhang, and Yuanchun Li. Parathinker: Native parallel thinking as a new paradigm to scale llm test-time compute. *arXiv preprint arXiv:2509.04475*, 2025.

xAI. Grok 3 beta the age of reasoning agents. <https://x.ai/news/grok-3>, February 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

Xinyu Yang, Yuwei An, Hongyi Liu, Tianqi Chen, and Beidi Chen. Multiverse: Your language models secretly decide how to parallelize and merge generation. *arXiv preprint arXiv:2506.09991*, 2025b.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822, 2023.

Yixin Ye, Yang Xiao, Tiantian Mi, and Pengfei Liu. Aime-preview: A rigorous and immediate evaluation framework for advanced mathematical reasoning. <https://github.com/GAIR-NLP/AIME-Preview>, 2025. GitHub repository.

Yiyao Yu, Yuxiang Zhang, Dongdong Zhang, Xiao Liang, Hengyuan Zhang, Xingxing Zhang, Ziyi Yang, Mahmoud Khademi, Hany Awadalla, Junjie Wang, et al. Chain-of-reasoning: Towards unified mathematical reasoning in large language models via a multi-paradigm perspective. *arXiv preprint arXiv:2501.11110*, 2025.

Di Zhang, Xiaoshui Huang, Dongzhan Zhou, Yuqiang Li, and Wanli Ouyang. Accessing gpt-4 level mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *arXiv preprint arXiv:2406.07394*, 2024.

Siyan Zhao, Devaansh Gupta, Qinling Zheng, and Aditya Grover. d1: Scaling reasoning in diffusion large language models via reinforcement learning. *arXiv preprint arXiv:2504.12216*, 2025.

Tong Zheng, Lichang Chen, Simeng Han, R Thomas McCoy, and Heng Huang. Learning to reason via mixture-of-thought for logical reasoning. *arXiv preprint arXiv:2505.15817*, 2025a.

Tong Zheng, Hongming Zhang, Wenhao Yu, Xiaoyang Wang, Xinyu Yang, Runpeng Dai, Rui Liu, Huiwen Bao, Chengsong Huang, Heng Huang, et al. Parallel-r1: Towards parallel thinking via reinforcement learning. *arXiv preprint arXiv:2509.07980*, 2025b.

A APPENDIX

In this appendix, we provide details omitted from the main text.

A.1 ALGORITHM: SSFT IMPLEMENTATION

Algorithm 1 Set Supervised Fine-tuning (SSFT)

Require: • π_θ : base model
 • N: Number of global forking tokens $\{g^{(i)}\}_{i=1}^N$
 • \mathcal{D} : Dataset with (at most) M reasoning traces per question
 • B: Global batch size
 • T_L : The first T_L number of tokens to match in matching cost (Equation 7).

Ensure: Output π_θ

- 1: **for** each training step **do**
- 2: **for** $k = 1, \dots, B$ **do**
- 3: Sample an input prompt and the corresponding reasoning traces $(\mathbf{x}_k, \{\mathbf{r}_k^{(j)}\}_{j=1}^M) \sim \mathcal{D}$
- 4: Initialize cost matrix $C \in \mathbb{R}^{N \times M}$
- 5: **for** $i = 1, \dots, N$ **do**
- 6: **for** $j = 1, \dots, M$ **do**
- 7: Compute the matching cost between $g^{(i)}$ and $\mathbf{r}_k^{(j)}$ by Equation 7.
- 8:
$$\mathcal{L}_{\text{matching}}(g^{(i)}, \mathbf{r}_k^{(j)}) = -\text{sg} \left(\frac{1}{T_L} \sum_{t=1}^{T_L} \log \pi_\theta \left(\mathbf{r}_{k,t}^{(j)} | \mathbf{x}_k, g^{(i)}, \mathbf{r}_{k,<t}^{(j)} \right) \right) \quad (7)$$
- 9: Store the matching cost in C.
- 10:
$$C(i, j) = \mathcal{L}_{\text{matching}}(g^{(i)}, \mathbf{r}_k^{(j)}) \quad (8)$$
- 11: **end for**
- 12: **end for**
- 13: Compute optimal matching $\hat{\sigma}_k$ between $\{g^{(i)}\}_{i=1}^N$ and $\{\mathbf{r}_k^{(j)}\}_{j=1}^M$. Hungarian algorithm (Kuhn, 1955) can be applied to C to efficiently compute Equation 9 (Equation 3).
- 14:
$$\hat{\sigma}_k = \arg \min_{\sigma \in \mathfrak{S}_P} \sum_{j=1}^M C(\sigma(j), j) \quad (9)$$
- 15: **end for**
- 16: Compute the empirical *set language modeling loss* (Equation 10):
- 17:
$$\mathcal{L}_{\text{Hungarian}}(\theta) = -\frac{1}{B} \sum_{k=1}^B \left[\sum_{j=1}^M \sum_{t=1}^{T_r} \log \pi_\theta \left(\mathbf{r}_{k,t}^{(j)} | \mathbf{x}_k, g^{(\hat{\sigma}_k(j))}, \mathbf{r}_{k,<t}^{(j)} \right) \right] \quad (10)$$
- 18: Update model parameters θ using gradients $\nabla_\theta \mathcal{L}_{\text{Hungarian}}(\theta)$
- 19: **end for**

Algorithm 1 presents the core SSFT implementation with optimal bipartite matching. **In practice, the nested-loop computation used to populate C is fully vectorized and can be executed in a single forward pass.** This does not blow up VRAM because (i) we do not store activations for these cost evaluations (no backprop through matching costs), and (ii) We only need to use the first $T_L < T_r$ NTP losses to compute the matching cost, as the NTP loss over the first few thousand tokens can already differentiate many reasoning traces in terms of their modes. Nevertheless, our code also supports matching over the full T_r tokens by chunking the computation into a few batches, so this step does not become a VRAM bottleneck. Fine-tuning on 1k questions with 4 traces each, SSFT (optimal matching) took 6.5 h for 6 epochs, compared to 6.1 h for standard SFT, adding only 6.6% overhead.

Algorithm 2 Queue-based Distributed SSFT with variable number of traces for each question

Require: • π_θ : base model
• N: Number of global forking tokens $\{g^{(i)}\}_{i=1}^N$
• \mathcal{D} : Dataset with (at most) M reasoning traces per question, variable m number of traces per question
• B: Global batch size
• T_L : The first T_L number of tokens to match in matching cost (Equation 7).
• b : Original per-device global batch size ($b \geq M$), This is “micro batch size*original grad accumulation steps”

Ensure: Output π_θ

1: **for** each epoch **do**

2: Initialize Queue Q for storing a sequence of $(\mathbf{x}_k, \{\mathbf{r}_k^{(j)}\}_{j=1}^m)$ where m is a variable number that differs between input questions and different processes (GPUs)

3: Initialize Queue q for storing a sequence of sizes of sets m .

4: **for** every $(\mathbf{x}_k, \{\mathbf{r}_k^{(j)}\}_{j=1}^m) \in \mathcal{D}$ **do**

5: $Q \leftarrow Q.enqueue((\mathbf{x}_k, \{\mathbf{r}_k^{(j)}\}_{j=1}^m))$

6: $q \leftarrow q.enqueue(m)$

7: $all_q_list = All-gather(q)$

8: Initialize list $temp_batch$

9: **while** All processes have at least b sequences based on all_q_list **do**

10: **while** $temp_batch$ does not have at least b sequences **do**

11: $temp_batch \leftarrow temp_batch.append(Q.dequeue())$

12: $q.dequeue()$

13: **end while**

14: compute the the maximum per_device global batch size b_{max} currently in all processes using all_q_list (inferred, no collective operation)

15: Pad $temp_batch$ to size b_{max} by appending “pad sequences” as needed.

16: Update all entries in all_q_list based on inferred usage

17: Perform one SSFT training step, SSFT ($\pi_\theta, temp_batch, b_{max}$)

18: **end while**

19: **end for**

20: **end for**

The primary VRAM bottleneck in SSFT remains the backpropagation Step 10, regardless of whether we use optimal or random matching, because the effective batch size scales with M . To mitigate this, we split the backward pass into several gradient-accumulation steps. Although our experiments use the same number of reasoning traces per question, we also support variable number of targets using our queue-based batching in Algorithm 2. The complication arises when using distributed training with a variable-sized batch, as different processes require the same per-device batch size to perform collective operations. This is mitigated by padding with “PAD sequences” to align batch sizes. Our implementation minimizes the number of “PAD sequences” by storing a variable number of targets in a queue and dequeuing multiple items to form a per-device global batch, so smaller batches can be stitched together instead of always being padded.

A.2 LEARNED MATCHINGS BETWEEN GLOBAL FORKING TOKENS AND TRACES

Initially, any of the bipartite matching configuration $\sigma_k \in \mathfrak{S}_P$ can be computed as optimal, as the reserved global forking tokens $\{g^{(i)}\}_{i=1}^N$ have no specific correlations with these traces $\{r^{(j)}\}_{j=1}^M$. This can be observed in Figures 7 and 8 that $c(\sigma_k)$, the count of configuration σ_k being optimal during training, uniformly increases on all configuration indexes. However, as training goes with SSFT, we notice only a subset of \mathfrak{S}_P accumulates mass. This indicates there are some unique correlations learned between $\{g^{(i)}\}_{i=1}^N$ and $\{r^{(j)}\}_{j=1}^M$. This subset is denoted as $\mathfrak{S}_p = \{\sigma_k\}_{k=1}^p$, and we call the unique edges in \mathfrak{S}_p as *learned matchings*.

A.2.1 HOW TO CHOOSE THE GLOBAL FORKING TOKEN FOR PASS@1

To find \mathfrak{S}_p , we can simply track which configurations σ still accumulate mass in the last epoch. Then we can connect all the unique edges in \mathfrak{S}_p to visualize learned matchings. However, multiple global forking tokens may share the maximum number of connected edges in the learned matchings. To break the tie, we treat the counts as edge weights and select $g^{(i)}$ with the largest weighted degree. We provide this implementation in our code.

A.3 TRAINING DETAILS

A.3.1 TRAINING DATASETS

32B experiments with questions from s1(Main): We explain the process of generating our training dataset for experiments in Table 1, Figure 4, Figure 5, Figure 6, Figure 7. First, we use the 1000 questions from s1 (Muennighoff et al., 2025) and populate a pool of reasoning traces by distilling from GPT-OSS-120B-high reasoning, GPT-OSS-120B-medium reasoning, DeepSeek R1, Gemini Flash2.0 Thinking, and Claude4/4.1. We use temperature 1.0, maximum length of 32768, and sets high reasoning effort unless specified. We generate two traces per teacher model. We use Claude3.5 Sonnet to extract the answer from the distilled solutions and compare with the ground-truth answer. The correctness of these distilled traces are shown in 3.

For s1k-4mixed-reasoning dataset, we sample 4 traces per question from this pool, so the dataset consists of 1000 questions, each paired with 4 reasoning traces. This dataset was used to fine-tune **SSFT-32B**, **SSFT-32B (random σ)**, and **SFT-mixed-distill-32B**.

For fine-tuning **SFT-OSS-distill-32B** model, we only use the 1000 traces from “Run1” GPT-OSS-120B with high reasoning effort.

For obtaining the visualizations in Figure 7 and Figure 8, we fine-tune models using the same teacher models for all 1000 questions. we always choose the 4 traces from “Run1” of GPT-OSS-120B-high, GPT-OSS-120B-medium, DeepSeek R1, and Gemini Flash2.0 Thinking. As mentioned in Section 3.3 and Figure 7, we fine-tune under 3 bipartite matching hyperparameters to conduct this case study.

A.3.2 TRAINING HYPERPARAMETERS

For consistency, we use Qwen2.5-32B-Instruct (Yang et al., 2025a) as the base model for all of our 32B experiments. We use standard fine-tuning hyperparameters: we train for 6 epochs with a global batch size of 32, which is derived from 4 gradient accumulation steps and distributed training with 8 GPUS ($4 \times 8 = 32$). This results in 756 gradient steps. The maximum sequence length is set to

	GPT-OSS-120B-high	GPT-OSS-120B-medium	DeepSeekR1	Gemini Flash2.0 Thinking	Claude Opus4/4.1
Run1	796/1000	769/1000	620/1000	538/1000	656/1000
Run2	785/1000	753/1000	641/1000	545/1000	647/1000

Table 3: The number of correct reasoning traces distilled for the 1,000 questions in s1 by different teacher models. This evaluation is done by Sonnet comparing the predictions and the ground-truth answers. We see that GPT-OSS has the highest accuracy for s1 dataset.

32,768. We train with bfloat16 with a learning rate of $1e-5$ warmed up linearly for 5% and then decayed to 0 using a cosine schedule. We choose AdamW optimizer (Loshchilov & Hutter, 2017) with $\beta_1 = 0.9$, $\beta_2 = 0.95$, and weight decay $1e-4$. We only backpropagates the completion loss, which is the loss on reasoning traces and the answers. Fine-tuning **SSFT-32B** plus loggings took 6.5 hours on 8 NVIDIA B200 GPUs using PyTorch FSDP, Liger Kernel (Hsu et al., 2024) for fused cross entropy loss, and FlashAttention-2 (Dao, 2023) for fused attention computation. Fine-tuning **SSFT-32B (random σ)** took 6.3 hours, and Fine-tuning **SFT-mixed-distill-32B** took 6.1 hours. Even our baseline **SFT-OSS-distill-32B** with only one trace per question, and our attempt to reproduce s1.1 took 1.66 hours, which is longer than the time reported by Muenninghoff et al. (2025). This is due to using 8 GPUs instead of 16 GPUs, hardware and package differences. When training with s1k-4mixed-reasoning, we added one extra epoch from 5 epochs to 6 epochs, since we have 4x reasoning traces, but we did not linearly increase the number of epochs, as these traces can be similar, and the number of distinct questions is still 1,000. Overall, we made sure all of our models are fine-tuned with consistent hyperparameters.

A.3.3 VISUALIZATION OF SSFT TRAINING DYNAMICS

Figure 9 shows the standard training dynamics of SSFT with optimal bipartite matching. The resulting model is SSFT-32B. The loss plotted here is Equation 5.



Figure 9: Training dynamics of SSFT-32B on s1k-4mixed-reasoning

Figure 10 shows the evolution of bipartite matching during SSFT. Figures 10a and 10b show that the gap between optimal bipartite matching cost and non-optimal bipartite matching cost under other σ keeps widening during training. This means that these reasoning traces are indeed starting to match unique global forking tokens. Even though SSFT effectively optimizes a non-stationary objective which depends on model parameters θ , the widening gap shows the inner discrete optimization is converging as training goes. Figure 10c also confirms that the model learned some unique correlations between $\{g^{(i)}\}_{i=1}^6$ and $\{r^{(j)}\}_{j=1}^4$, as only a subset of matchings are still computed as optimal. Compared to Figure 7, we see more σ accumulating mass towards the end. This is due to having more mixed diverse reasoning traces, so the model learned more intricate associations between these global forking tokens and truly diverse reasoning traces.

A.3.4 ABLATION STUDY TRAINING DETAILS (REMOVING HIGH QUALITY SMALL DATASET)

For this ablation study, we choose Open-R1-Math220k default split, which has 93,000 math questions and 2 \sim 4 traces. Since Qwen2.5-Math-7B is a widely fine-tuned model using this dataset, we also choose it as our base model. We train for 3 epochs using 8 A100 GPUs, which took around 4

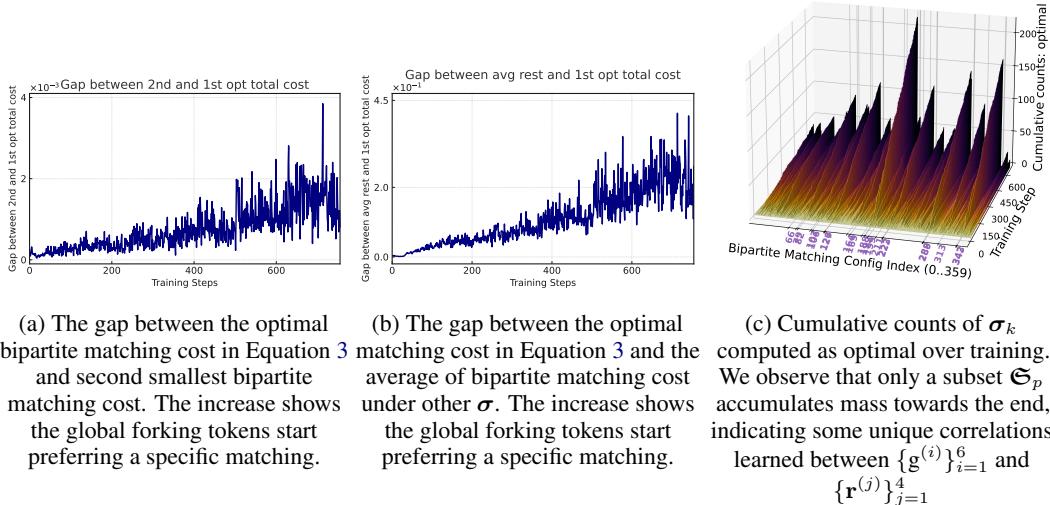


Figure 10: Dynamics of bipartite matching during the fine-tuning of SSFT-32B on s1k-4mixed-reasoning.

days. Our hyperparameters are mostly consistent with the recommended hyperparameters by Face (2025). We fine-tune both SSFT-OpenR1-93k-7B and SFT-OpenR1-93k-7B with a maximum length of 32768, learning rate of $4.0e-05$ warmed up linearly for 3% and decayed to 0 following cosine schedule, 8 gradient accumulation steps. For our SSFT method, we reserve $N = 4$ global forking tokens, and use the first 1000 tokens for matching. Again, only the completion loss is used for optimizing the model parameters.

A.4 EXAMPLES OF SSFT-32B REASONING TRACES INITIATED BY DIFFERENT GLOBAL FORKING TOKENS

We randomly picked a set of six parallel generations from those where SSFT-32B correctly answered a problem from AIME 2025 under all six global forking tokens. This is question 11.

Problem 11 (AIME25)

A piecewise linear periodic function is defined by

$$f(x) = \begin{cases} x, & \text{if } x \in [-1, 1), \\ 2 - x, & \text{if } x \in [1, 3), \end{cases} \quad \text{and} \quad f(x + 4) = f(x) \text{ for all real } x.$$

The graph of $f(x)$ has a sawtooth pattern. The parabola $x = 34y^2$ intersects the graph of $f(x)$ at finitely many points. The sum of the y -coordinates of these intersection points can be expressed in the form

$$\frac{a + b\sqrt{c}}{d},$$

where a, b, c, d are positive integers, $\gcd(a, b, d) = 1$, and c is squarefree. Find $a + b + c + d$.

Here are the 6 parallel generations by SSFT-32B. We see distinct reasoning traces produced by prompting with a unique $g^{(i)} \in \{g^{(i)}\}_{i=1}^6$. The differences are highlighted in red comments.

<think1> — Periodization + Quadratic Families + Vieta/Telescoping

```
<| im_start |>assistant
<| im_start |>
<think1>
```

“write $x = t + 4k$ with $t \in [-1, 3]$ ” “Case 1 ($t \in [-1, 1]$, $y = t$) ... Case 2 ($t \in [1, 3]$, $y = 2 - t$)” $34t^2 - t - 4k = 0$ and $t = \frac{1 \pm \sqrt{1 + 544k}}{68}$ “let $u = 2 - t$ ” and $u = \frac{-1 \pm \sqrt{273 + 544k}}{68}$ Thus $u = \frac{-1 + 5\sqrt{185}}{68}$ “telescopes to rational constant $9/34$ ” (*Comment:* periodize $x = t + 4k$, solve two quadratic families, use Vieta pair-sums; edge-period gives the irrational y)

</think1>
<| im_start |>answer
Answer: 259

<think2> — Nearest-Multiple-of-4 Reduction

<| im_start |>assistant
<| im_start |>
<think2>
Start from $y = f(34y^2)$ write $34y^2 = 4k + r$ with $r \in [-1, 3]$ (choose $k = \lfloor \frac{34y^2+1}{4} \rfloor$) Branch A ($r \in [-1, 1]$): $y = r \Rightarrow 34y^2 - y - 4k = 0$ Branch B ($r \in [1, 3]$): $y = 2 - r \Rightarrow 34y^2 + y - (2 + 4k) = 0$ Vieta: pair-sums $+\frac{1}{34}$ and $-\frac{1}{34}$ cancel over full periods only the edge/partial period remains ($k = 8$, descending branch)
 $y = \frac{-1 + \sqrt{273 + 544 \cdot 8}}{68} = \frac{-1 + 5\sqrt{185}}{68}$
 (*Comment:* reduce to $r \in [-1, 3]$ via nearest multiple of 4; full periods cancel by $\pm \frac{1}{34}$ pair-sums; only the boundary $k = 8$ root contributes)

</think2>
<| im_start |>answer
Answer: 259

<think3> — Paper-Style Sections; Explicit Pair-Sums and Totals

<| im_start |>assistant
<| im_start |>
<think3>
“Now we need sum of all y -coordinates” $\sum_{n=0}^8 y_{1,n} + \sum_{n=0}^7 y_{C,n}$ $y = \frac{-1 + \sqrt{4625}}{68}$
 “Therefore $S_2 = \frac{5\sqrt{185} - 17}{68}$,” “Total sum S ...” (*Comment:* formal write-up; separates sums by family, explicitly displays the boundary $y = \frac{-1 + \sqrt{4625}}{68}$, then totals)

</think3>
<| im_start |>answer
Answer: 259

<think4> — Set/System Framing Before Algebra

<| im_start |>assistant
<| im_start |>
<think4>
“the set $\{(x, f(x))\}$... Intersecting it with the parabola $x = 34y^2$ requires solving $y = f(x)$, $x = 34y^2$.” (then proceeds to the same two-case algebra)

$u_+ = \frac{-1 + \sqrt{273 + 544 \cdot 8}}{68} = \frac{-1 + 5\sqrt{185}}{68}$ (*[Comment:]* definition-first framing; after setting up the system, it reaches the same boundary u_+ value on the partial period)
</think4>
<| im_start |>answer
Answer: 259

<think5> — Outline & Symmetry: Full Periods Cancel, Last Period Contributes

<| im_start |>assistant
<| im_start |>
<think5>
“the sum is zero ... That’s a key insight.” “for each full period ($k = 0$ to $k = 7$), the sum ... is zero ... now consider the partial period $k = 8$ ” Compute $u_{pos}(k=8) = \frac{-1 + \sqrt{4625}}{68}$ (term-by-term cancellation examples for $k = 0, 1$) (*[Comment:]* symmetry explains full-period cancellation; the lone boundary contribution at $k = 8$ supplies the irrational term)
</think5>
<| im_start |>answer
Answer: 259

<think6> — Stream-of-Consciousness Scratchwork

<| im_start |>assistant
<| im_start |>
<think6>
“Okay, so I need to solve this problem ... Let me try to break this down step by step.” ...
... “Wait, actually, no. Wait,” $y_{A,8} = \frac{-1 + \sqrt{4625}}{68}$ (*[Comment:]* exploratory self-checks; reaches the same boundary y on the partial period)
</think6>
<| im_start |>answer
Answer: 259

A.5 ALTERNATIVE IMPLEMENTATION OF TRAINING WITH MULTI-TARGET SEQUENCES

Figure 11 shows an alternative implementation for training with multiple target parallel sequences. Instead of flattening multi-target data along the batch dimension, we can also concatenate parallel reasoning traces along the sequence dimension, modify the causal attention matrix, and position ids for positional embeddings. While easy to implement, this version extends the sequence length and cannot perform gradient accumulation along the batch dimension to maintain the same VRAM as training with a single reasoning trace. Since most fine-tuning already sets the micro-batch size to 1 due to the increased number of reasoning tokens already present in a single reasoning path, this implementation also cannot further shrink the micro-batch size to accommodate the extended context length. In addition, the code that uses the initial flash-attention implementation, which requires causal attention matrix, is not compatible with this setup. This implementation is also less flexible in maintaining a consistent number of sequences per global batch.

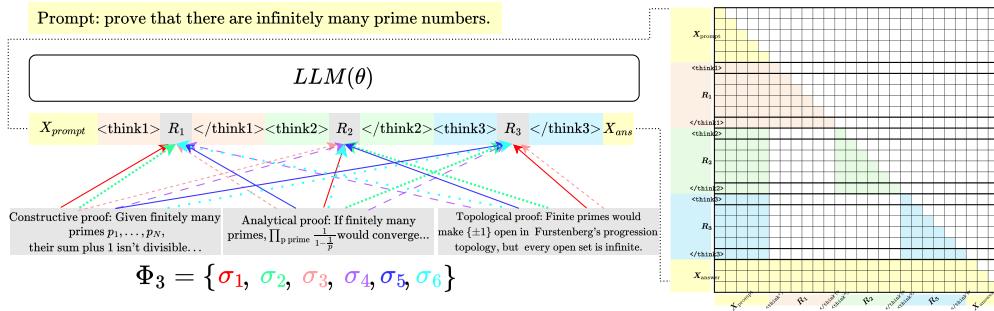


Figure 11: An illustration of training with multiple parallel targets along the sequence dimension.