AgentArk: Distilling Multi-Agent Intelligence into a Single LLM Agent

Yinyi Luo[1]*, **Yiqiao Jin**[3], **Weichen Yu**[1], **Mengqi Zhang**[2], **Srijan Kumar**[3], **Xiaoxiao Li**[5], **Weijie Xu**[4]†, **Xin Chen**[4], **Jindong Wang**[2]†

[1]Carnegie Mellon University [2]William & Mary [3]Georgia Institute of Technology
[4]Amazon [5]University of British Columbia

Abstract

While large language model (LLM) multi-agent systems achieve superior reasoning performance through iterative debate, practical deployment is limited by their high computational cost and error propagation. This paper proposes **AgentArk**, a novel framework to distill multi-agent dynamics into the weights of a *single* model, effectively transforming explicit test-time interactions into implicit model capabilities. This equips a single agent with the intelligence of multi-agent systems while remaining computationally efficient. Specifically, we investigate three hierarchical distillation strategies across various models, tasks, scaling, and scenarios: reasoning-enhanced fine-tuning; trajectory-based augmentation; and process-aware distillation. By shifting the burden of computation from inference to training, the distilled models preserve the efficiency of one agent while exhibiting strong reasoning and self-correction performance of multiple agents. They further demonstrate enhanced robustness and generalization across diverse reasoning tasks. We hope this work can shed light on future research on efficient and robust multi-agent development. Our code is available at `https://github.com/AIFrontierLab/AgentArk`.

1 Introduction

Multi-agent Systems (MAS), where multiple models interact through debate [11, 17, 13], critique [25, 67], and consensus [4], have demonstrated remarkable success in complex reasoning tasks [14]. By structuring reasoning as a multi-turn and multi-role dialogue, MAS can explore diverse hypotheses, uncover logical errors [59], and iteratively refine solutions [51]. However, this collaborative power is a double-edged sword that introduces systemic risks including *computational overhead* and *vulnerability amplification*. First, the reliance on multi-role and multi-turn dialogue causes inference latency and computational overhead to grow rapidly [58, 23]. In densely connected networks, computation can grow quadratically with the number of agents, making MAS prohibitively expensive for real-time [24, 19]. Second, while MAS can correct errors, they can also escalate them. In high-density interactions, individual biases or hallucinations can propagate and amplify across the group, leading to collective failures in robustness and safety [16, 42].

These challenges naturally raise a fundamental question: *Can a single model internalize the reasoning benefits of MAS without their high inference-time cost and collaborative vulnerabilities?*

The collective power of MAS, i.e., *inference-time compute*, suggests that the gains can possibly be "shifted forward"—internalized by a single model during offline learning [54, 7, 38]. While prior work shows that single models can absorb certain MAS reasoning benefits [29, 28, 5, 76], they are often limited to imitating final answers [15] or exhibiting shallow interaction traces [27, 35], failing

*Contact author: `yinyil@andrew.cmu.edu`.
†Corresponding authors: `jdw@wm.edu` and `weijiexu@amazon.com`.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

to reproduce the core *iterative conflict-and-refinement dynamics* that underlie MAS reasoning [30]. Notably, recent evidence suggests that the structural design of MAS may play a secondary role in observed gains. [24] demonstrate that removing or perturbing explicit agent structures leads to only marginal performance degradation, revealing the essential contribution of MAS lies in the *reasoning dynamics* they induce, rather than in the interaction schema itself. Similarly, [22] systematically studies orchestration strategies under controlled benchmarks, showing that performance is largely driven by induced reasoning behaviors rather than agent topologies or protocols.

Our pilot study showed that fine-tuning only on final answers can lead to overfitting on task-specific mapping with limited generalization (Appendix G.1). Therefore, a single agent should learn to internalize the reasoning dynamics, allowing for generating, evaluating, and refinement with one forward pass. This perspective also resonates with human cognition: individuals can internalize group reasoning strategies and reproduce collective wisdom independently [50, 41, 9]. In this paper, we present a comprehensive investigation of multi-agent distillation for reasoning tasks. Particularly, we propose **AgentArk** (Figure 1), a general distillation framework that transfers MAS reasoning dynamics into a single model, without relying on handcrafted interaction or task-specific supervision.[3]

AgentArk integrates three progressively deeper distillation strategies: 1) Reasoning-Enhanced Outcome–Based Supervision *(R-SFT)*: Training the model on the final consensus reached by the agent group while leveraging reasoning trajectories; 2) Reasoning Trajectory-based Data Augmentation *(DA)*: Distilling the diverse reasoning chains derived from collective interaction, allowing the model to learn various logical strategies and ways of thinking; and 3) Process-Aware Distillation *(PAD)*: Using process reward model to train the agent to internalize the critique-and-revision dynamics, enabling a single agent to emulate the dialectical reasoning of multi-agent debates within a single forward pass. We conduct extensive experiments by varying LLM backbones, teacher-student roles, datasets and various tasks, scaling, and evaluation settings. The following are our key findings:

1. **AgentArk enables a single agent to acquire multi-agent reasoning ability.** Our extensive experiments show that all three reasoning-centric distillation methods can boost the performance of single agents. Combination of approaches can achieve further improvement. (§4.2)
2. **PRM capacity matters more than student model size, while student capacity bounds multi-agent gains.** High-capacity PRMs enable strong improvements even for small students, whereas weak PRMs limit gains. Scaling teacher agents mainly benefits larger students, with diminishing returns for smaller ones. (§4.2, §4.3)
3. **Reasoning quality outweighs quantity.** Simply adding more reasoning trajectories does not improve performance, while PAD's high-signal process supervision yields stable gains. (§4.3)
4. **Process-aware distillation improves reasoning behavior, not just accuracy.** PAD models exhibit better step decomposition, self-checking, and error correction than RSFT and RA. (§4.4)
5. **AgentArk improves generalization and robustness.** Distilled models transfer reliably to unseen and robustness benchmarks. (§4.5)
6. **AgentArk extends to multimodal and heterogeneous LLMs and other MAS algorithms.** (§4.6)

Contributions. (1) To our best knowledge, AgentArk is the *first* comprehensive framework to explore various strategies for MAS distillation. (2) We construct a scalable distillation data generation pipeline and framework agnostic to MAS strategies, which will be released to facilitate future research on reasoning distillation. (3) We perform extensive evaluation of MAS distillation from various perspectives, providing insights for future research.

2 Related Work

Multi-Agent Systems. MASs have emerged as an effective paradigm for enhancing LLM reasoning by enabling multiple agents to interact [68, 4, 13, 11, 63]. By organizing reasoning as explicit multi-turn interactions, they can explore diverse solution paths, detect errors, and iteratively refine predictions, leading to strong performance on complex reasoning tasks [59, 61, 62]. However, MAS rely on inference-time coordination among multiple agents, which incurs substantial computational cost and latency [48, 39, 56, 57]. In addition, agent roles, interaction protocols, and evaluation criteria are typically designed for specific tasks, limiting the applicability in resource-constrained or real-time settings. Beyond efficiency, recent work [24] examined the structural sensitivity of MAS, showing that performance is often robust to substantial simplifications or perturbations of agent structures.

[3] We only focus on reasoning tasks; tool use [45] and memory management [70] are for future work. We mainly consider the popular homogeneous setting where all agents share the same LLM backbone [13, 3]. Heterogeneous results are in §4.6.

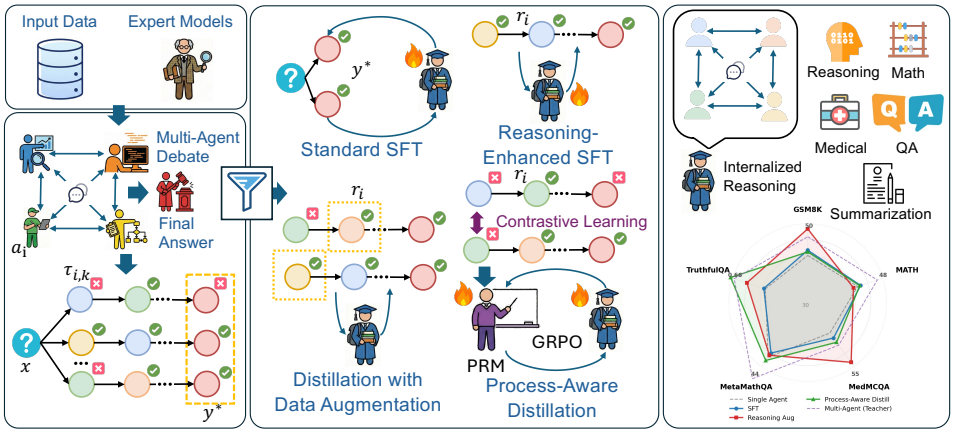

Figure 1: **Overview of AgentArk**: (1) **Data Generation Through Multi-Agent Debate** to produce diverse reasoning trajectories; (2) **Knowledge Extraction** to filters for high-quality corrective traces; and (3) **Distillation** utilizing Standard SFT, Reasoning-enhanced SFT, Distillation with Data Augmentation, and Process-Aware Distillation (PRM + GRPO).

Distillation of Multi-Agent Reasoning. To avoid the above issues, recent work has explored distilling multi-agent reasoning into a single model [36, 72, 53, 44]. Early approaches supervise student models using the final outputs of agent groups or simplified interaction traces [37, 21]. More recent methods transfer richer supervision signals, including graph-based interaction modeling [5], skill selection frameworks [29], debate-derived preference supervision [76] and end-to-end agentic reinforcement learning [28]. Despite their effectiveness, they often depend on task-specific agent designs, predefined interaction structures, or environment-dependent reward functions, restricting generalization across tasks and limit scalability [52]. In contrast, AgentArk abstracts away agent roles and interaction structures by distilling interaction reasoning processes at the process level, enabling a single student model to generalize across tasks and agent configurations without task-specific design.

3 Method

As shown in Figure 1, AgentArk consists of three phases: (1) **Data Generation via Multi-Agent Interaction**, where teacher models generate diverse reasoning trajectories. Here, we leverage the iterative reflection and error-correction trajectories inherent in LLM debates [11]. (2) **Knowledge Extraction**, where successful and corrective traces are generated and filtered; and (3) **Hierarchical Distillation**, where the student is trained via supervised learning and process-level RL. We explore three paradigms: (1) Reasoning-enhanced Supervised Fine-Tuning (RSFT), (2) Data Augmentation via Diverse Extraction (DA), and (3) Process-Aware Distillation using PRM and GRPO.

3.1 Data Generation and Knowledge Extraction

We adopt debate [31] as the MAS mechanism for data generation with rich reasoning dynamics. Interactions where individuals reflect, revise, and converge have long been recognized as hallmarks of intelligent behavior, and debate naturally surfaces processes such as self-correction, error discovery, and cross-examination. Prior work on debate [11] demonstrates that structured disagreement can enhance both factual accuracy and reasoning quality. Results on other MAS algorithms are in §4.6.

Debate Setup. For each input x, we initialize n debating agents $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ sharing the same LLM. The agents engage in a K-round interaction. In each round, each agent a_i generates a reasoning trace $\tau_{i,k}$ based on the problem x and the previous traces of its peers $\{\tau_{j,k-1}\}_{j \neq i}$. This setup encourages agents to identify logical inconsistencies in others' arguments and iteratively refine their own reasoning traces. This process continues until K rounds is reached or a consensus emerges. The result is a comprehensive debate log $\mathcal{L}_x$ containing a set of final reasoning traces $\{r_1, r_2, \ldots, r_n\}$ and their corresponding answers.

Correctness-First Trajectory Selection. While traditional distillation typically utilizes error-free paths, we prioritize **corrective trajectories**, reasoning traces where an agent initially proposed an incorrect step but recognize successfully pivoted to a correct answer y^* after receiving critiques. For

each task, we extract 1) the final consensus answer y^*, verified against ground-truth labels; 2) the intermediate reasoning traces $\{r_i\}$ that successfully lead to y^*.

Knowledge Extraction. We extract multiple answer-consistent yet structurally diverse reasoning trajectories from debates, capturing both correct solutions and explicit self-correction behaviors.

3.2 Distillation Methods

We explore three strategies to supervise the student model π_θ, aiming to distill not only correct solutions but also concise and structured reasoning [75].

3.2.1 Reasoning-Enhanced SFT

Reasoning-enhanced SFT not only takes the final answers, but also the reasoning traces as supervision. The student model is trained to maximize the likelihood of the raw multi-agent reasoning traces. Given an input x and a successful trace $r = (r_1, \ldots, r_n)$ to final answer y^*, the objective is:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{(x,r,y^*)\sim\mathcal{D}}\mathcal{L}_{\text{res}} + \mathcal{L}_{\text{ans}}, \tag{1}$$

where

$$\mathcal{L}_{\text{res}} = \sum_{t=1}^{|r|} \log p_\theta(r_t \mid r_{<t}, x), \qquad \text{(Reasoning)}$$

$$\mathcal{L}_{\text{ans}} = \log p_\theta(y^* \mid r, x). \qquad \text{(Answer)}$$

The objective consists of two components: a reasoning loss $\mathcal{L}_{\text{res}}$, which optimizes the model's ability to generate a coherent sequence of intermediate rationales r, and an answer loss $\mathcal{L}_{\text{ans}}$, which ensures the final prediction y^* is grounded in both the input context x and the preceding reasoning path. SFT assesses the student's capacity for basic imitation of the multi-agent reasoning style, testing whether it can maintain structural consistency across the entire generation sequence.

3.2.2 Distillation with Data Augmentation

To fully exploit the diversity and multi-perspective nature of the multi-agent debate, we implement a *Correctness-First Diverse Extraction* strategy. For each problem x, we filter the agents to define a set of successful contributors $\mathcal{A}_{\text{correct}} \subseteq \mathcal{A}$. We then utilize a high-capacity teacher LLM as a "distiller" to parse the raw debate logs of these agents. The distiller is tasked with extracting $k \in \{1, 2, 3\}$ reasoning trajectories that are: (1) **Correct:** They must lead strictly to the ground-truth y^*. (2) **Diverse:** The teacher is prompted to select traces that employ distinct mathematical identities, different logical heuristics, or varied starting assumptions (details in Appendix A). The student is trained on this augmented set $\mathcal{D}_{\text{aug}}$ using:

$$\mathcal{L}_{\text{Aug}}(\theta) = -\frac{1}{k}\sum_{i=1}^{k}\sum_{t=1}^{T} \log p_\theta(y_t \mid y_{<t}, r_i, x). \tag{2}$$

This forces model to learn multiple valid paths to the same solution for robustness and generalization.

3.2.3 Process-Aware Distillation

The third method treats distillation as a reinforcement learning problem, using a Process Reward Model (PRM) [32, 55] for granular, step-level supervision that reinforces the intermediate logical transitions found in multi-agent debates. Specifically, we train a PRM R_ϕ to predict the probability that a reasoning step is correct. To ensure the PRM reliably captures the nuances of the debate, we use a two-stage curriculum: **Stage I: Feature Alignment (Backbone Frozen).** We initialize R_ϕ with the weights of the student model and freeze all layers except for the final one and the reward head. This prevents the loss of pre-trained linguistic features while the reward head learns to map existing representations to the correctness labels $z_t \in \{0, 1\}$. **Stage II: Full Specialization (Backbone Unfrozen).** We unfreeze the entire backbone for end-to-end fine-tuning. This allows the model to develop specialized internal attention patterns for detecting logical fallacies.

We design the Process Reward Model (PRM) to provide step-level supervision via a **contrastive reward objective** rather than standard binary cross-entropy. This encourages the model to assign higher rewards to reasoning steps that are more consistent with the multi-agent debate consensus, reflecting relative correctness rather than absolute labels (details in Appendix B.1).

Reinforcement Learning via GRPO. Finally, we fine-tune the student policy π_θ using Group Relative Policy Optimization (GRPO) [47]. Given an input $x \sim \mathcal{D}$, we sample a group of G reasoning outputs $o_1, \ldots, o_G$ from a fixed behavior policy π_{old}, which is a snapshot of the student policy before the current update. GRPO optimizes the policy by comparing the rewards of these outputs within the group, removing the need for a separate value function:

$$\mathcal{J}(\theta) = \mathbb{E}_{x\sim\mathcal{D},\{o_i\}\sim\pi_{\text{old}}} \left[\frac{1}{G} \sum_{i=1}^{G} \mathcal{L}_i(\theta) - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right] \tag{3}$$

where π_{ref} denotes a fixed reference policy used to regularize the update, β is the KL penalty coefficient, and $\mathcal{L}_i(\theta)$ is the clipped surrogate objective for output o_i, the surrogate objective for each output o_i is:

$$\mathcal{L}_i(\theta) = \min \left(\rho_i(\theta)\hat{A}_i, \text{clip}(\rho_i(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_i \right). \tag{4}$$

In this formulation, $\rho_i(\theta) = \frac{\pi_\theta(o_i|x)}{\pi_{\text{old}}(o_i|x)}$ denotes the probability ratio. The advantage $\hat{A}_i$ ensures **reward consistency** by normalizing the step-wise aggregate score $R_\phi(o_i)$ from the trained PRM: $\hat{A}_i = (R_\phi(o_i) - \mu_R)/\sigma_R$, where μ_R and σ_R are the mean and standard deviation of the PRM scores in group G sampled outputs.

4 Experiments

4.1 Experimental Setup

Models. We conduct evaluation across three major model families: (1) **Qwen 3** [64]: Qwen3-32B, Qwen3-8B, Qwen3-1.7B, and Qwen3-0.6B; (2) **Gemma 3** [49]: Gemma3-27B-it and Gemma-7B; and (3) **Llama 3** [12]: Llama3-8B-Instruct. Particularly, we perform distillation from larger models (Qwen3-32B, and Gemma3-27B-it) to smaller ones (Qwen3-8B, Qwen3-1.7B, Qwen3-0.6B, Llama3-8B, and Gemma-7B) for comprehensive evaluation. We additionally evaluate multimodal LLM, heterogeneous agents, and other MAS algorithms as a preliminary study in §4.6.

Datasets. We adopt diverse benchmarks: (1) MATH [32] and GSM8K [8] for *mathematical reasoning*; (2) MedMCQA [43] for *domain-specific knowledge*; (3) MetaMathQA (MMQA) [66] for augmented math; and (4) QASPER [10], HotpotQA [65] and QMSum [74] for *multi-hop and long-form reasoning*.

Comparison Methods. We mainly compare three distillation methods in §3: RSFT, DA, and PAD, with single-agent and the vanilla multi-agent debate as baselines. Our primary study are based on 5 agents following existing work [31] and §4.3 presents the scaling exploration. By varying teacher-student models and training-test datasets on different distillation methods, we conducted 120 experiments in total. Accuracy on the test dataset is adopted as the primary evaluation metric.

4.2 Main Results

Figure 2 presents the results by distilling from Qwen3-32B to different student models across various datasets and the right-hand-side of Figure 1 shows the average comparison results. Other results are in Appendix G.2. Overall speaking, AgentArk significantly improves the performance of a single agent by 4.8% and only slightly worse than the vanilla MAS. More insightful findings are as follows.

Performance across distillation methods. Figure 2 show the results tested on GSM8K and MedMCQA, respectively. (1) **ID vs. OOD.** While distillation shows improvement in both settings, the gain in ID is more significant than in the OOD setting (30% vs. 7% maximum improvement and 4-6% vs. 1-3% on average). This is expected as OOD is more challenging, suggesting more future work. (2) **Methods.** Different methods show varied performance across datasets and tasks. While RSFT and DA sometimes provide improvements, their gains are inconsistent and can fluctuate by dataset.

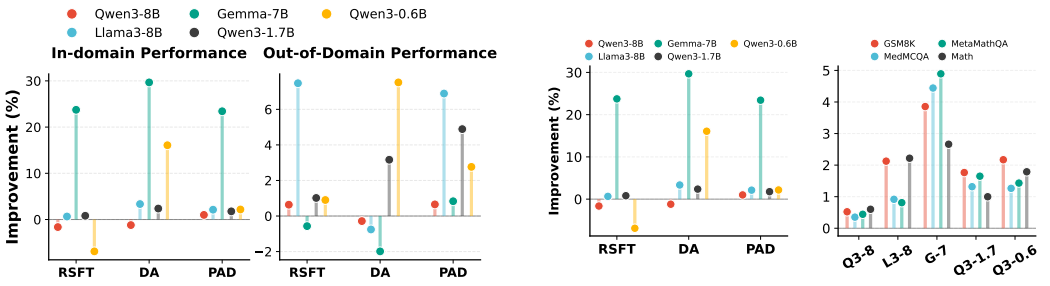

(a) In-domain (left) vs OOD (right) (b) By datasets (left) and models (right)

Figure 2: Distillation from Qwen3-32B to different student models.

In contrast, PAD consistently yields performance improvements, demonstrating its robustness and reliable transfer of reasoning capabilities. (3) **Compatibility.** Different distillation strategies are mutually compatible and can be composed to yield consistent gains (Appendix F).

Performance across student models. We evaluate the generalization of all three distillation strategies by fixing the teacher model and varying the student model family. The results are shown in Figure 2a and 2b (left). (1) **Same-family distillation (Qwen-3)**. When both teacher and student are drawn from the Qwen-3 family, distillation yields stable but relatively moderate improvements, with smaller variants (1.7B and 0.6B) consistently benefiting more than the 8B model, indicating that same-family distillation mainly alleviates capacity constraints rather than inducing substantial representational changes. (2) **Cross-family distillation.** When distillation is performed across different model families, we observe larger and more consistent performance gains, particularly for Gemma-7B and LLaMA-3-8B, indicating that heterogeneous architectures benefit more from transferred reasoning signals. (3) **Effect of PRM supervision.** Across both large and small student models, PRM-based distillation consistently improves performance on ID and OOD tasks, demonstrating that the learned gains are not limited to scale or data distribution. Notably, the improvements persist under distribution shift, suggesting that PRM supervision primarily transfers reasoning behaviors rather than merely improving surface-level alignment to training data (Detailed results in Appendix G.2). We further conduct detailed PRM ablation studies to analyze its effects (Appendix C).

Performance across datasets. For a fixed student model, distillation consistently improves performance across all benchmark datasets, but the magnitude of gains varies (Figure 2). MetaMathQA exhibits the largest improvements, followed by GSM8K, while Math shows moderate gains and MedMCQA benefits the least. This pattern suggests that the observed improvements are primarily driven by enhanced reasoning capabilities rather than dataset-specific overfitting. In particular, the strong gains on MetaMathQA and GSM8K indicate that these datasets contain reasoning-intensive problems, such as multi-step logical or arithmetic chains, which benefit most from transferred reasoning knowledge. By contrast, Math, with more formulaic or domain-specific tasks, shows moderate improvement, and MedMCQA, which relies heavily on specialized medical factual knowledge, benefits the least, implying that distillation contributes less when reasoning is minimal. Overall, this analysis highlights that our distillation strategies effectively enhance generalized reasoning skills, with larger impact on datasets that require complex reasoning rather than memorization.

4.3 Scaling and Data Dynamics

Scaling the Number of Agents. While the primary experiments employ 5 agents following existing work, we further scale to 10 and 20 agents to study how increasing teacher diversity and interaction complexity affects distillation. We perform evaluation by distilling from a fixed Qwen3-32B model into two different student models: Qwen3-8B and Qwen3-0.6B. The results are shown in Figure 3.

For the smaller student 0.6B, scaling > 5 does not yield additional benefits and, in some cases, leads to performance degradation. We attribute this to student's limited representational capacity: as the teacher ensemble becomes more diverse and produces more complex or longer reasoning trajectories, the student is unable to faithfully absorb and generalize this information. By contrast, the larger Qwen3-8B benefits modestly from scaling. Performance improves consistently as the number of agents increases. However, the incremental gains diminish at higher scales, indicating that once the student can effectively utilize the ensemble's reasoning, additional agents offer limited benefit. These results highlight that MAS distillation is bounded by student capacity: smaller models saturate

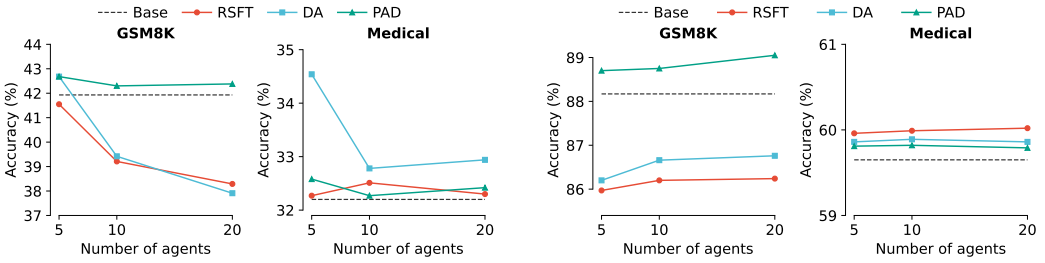

(a) 0.6B on GSM8K (b) 8B on GSM8K

Figure 3: Agent scale $(5, 10, 20)$ on distillation performance evaluated on GSM8K and MedMCQA.

quickly and may even struggle with overly diverse teacher signals, whereas larger models can better exploit richer supervision from multiple agents. Overall, scaling teachers is effective only when matched to the student's representational capacity (More results in Appendix D.2).

Data Quantity vs. Quality. The amount of training data is varied in Figure 4, showing that increasing training data does not lead to monotonic improvements. For both RSFT and DA, performance exhibits high variance as data scale grows: moderate data sizes can yield gains, while further scaling often leads to stagnation or even degradation, particularly for GSM8K and MedMCQA. In contrast, PAD demonstrates significantly more stable behavior across data scales. They consistently achieve competitive performance and avoid the sharp fluctuations observed with raw data scaling. This suggests that for capacity-limited students, reasoning quality rather than data volume is the primary bottleneck. Excessive or noisy supervision from large-scale MAS outputs can overwhelm the target, whereas PRM-guided distillation preserves high-signal reasoning trajectories for more reliable transfer.

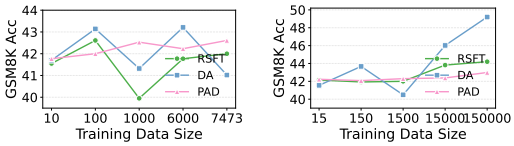

(a) GSM8K (b) MetaMathQA

Figure 4: Data scaling results (distillation from Qwen3-32B to Qwen3-0.6B).

Training Time. The offline training cost is shown in Appendix D.3, which is mild and does not affect the inference efficiency of a single agent.

4.4 Analysis of Reasoning Distillation

Comparison with On-Policy Self-Distillation. We compare with on-policy self-distillation (SD) [71], where the model improves itself using its own generated trajectories. Our PAD method outperforms SD (+0.99% vs. +1.62% on GSM8K). The results highlight a fundamental distinction in supervision signals. Specifically, SD relies on self-generated trajectories, which are constrained by the exploration capability of the student model; AgentArk leverages teacher-induced MAS reasoning trajectories, providing richer and more structured supervision that captures diverse reasoning behaviors beyond the student's own exploration. Overall, these results suggest that MAS distillation offers a more informative and structured learning signal than purely on-policy self-distillation, leading to more effective reasoning transfer.

Perplexity Analysis. We then conduct reasoning quality analysis by combining perplexity and LLM-based qualitative evaluation. For each dataset, we randomly sample 100 examples. We measure the perplexity ([2]; details in Appendix E.1) of distilled models on held-out samples from GSM8K, focusing on reasoning tokens using Qwen3-32B and 0.6B as the teacher and student models, respectively. Perplexity measures how well a model predicts the next token in a sequence; lower values indicate the model's reasoning steps are more predictable and coherent, aligning with our goal of producing structured reasoning trajectories. Table 1a shows that all distilled models achieve substantially lower reasoning perplexity than the single ones, indicating that distillation improves the predictability of the student's reasoning. This suggests that distillation not only helps produce more fluent outputs, but also encourages more structured and consistent reasoning trajectories.

Evaluation of Reasoning Quality. For more quantitative evaluation, we employ InternLM-2.5-20b-chat [20] as an automatic evaluator to score outputs along four dimensions: step decomposition

Table 1: Evaluation of reasoning quality and robustness.

(a) Quantitative evaluation of reasoning quality.

Metric	Single	RSFT	DA	PAD
Avg NLL (↓)	0.6529	0.4092	0.4449	0.5876
Perplexity (↓)	1.9211	1.6388	1.5603	1.7996
Step Decomposition (↑)	2.75	3.13	3.38	3.23
Intermediate Verification (↑)	2.41	3.48	4.04	4.07
Error Localization (↑)	1.97	2.19	2.91	2.78
Reasoning Coherence (↑)	1.88	2.25	3.07	3.96

(b) Robustness evaluation on TruthfulQA.

Metric	Single	PAD	DA	RSFT
BLEU (↑)	0.6034	0.6634	0.6353	0.6059
ROUGE-1 (↑)	0.6144	0.6659	0.6389	0.6157
ROUGE-2 (↑)	0.5704	0.6414	0.5961	0.5777
ROUGE-L (↑)	0.6132	0.6573	0.6401	0.6157

[18], intermediate verification [73], error localization [40], and overall reasoning coherence [26]. Step decomposition evaluates whether the model explicitly breaks problems into logical substeps; intermediate verification measures self-checking at each step; error localization assesses the model's ability to identify and correct mistakes; and overall reasoning coherence captures the consistency and logical flow of the solution. As shown in Table 1a, PAD achieves the highest scores in all metrics, indicating that it most effectively preserves explicit multi-step structure, self-checking behavior, and coherent reasoning flows. In contrast, DA shows moderate improvements, particularly in intermediate verification and error localization, suggesting that it captures surface-level reasoning structure without fully inheriting reflective reasoning behaviors. RSFT only outperforms the baseline, implying that direct reasoning-level supervision alone is insufficient.

Case Study. We further presented detailed comparisons of reasoning states in Appendix E. This example demonstrates that the distilled single-agent model acquires the multi-agent reasoning patterns: its reasoning is more structured, logically coherent, and self-consistent, producing the correct answer without the repeated self-corrections seen in the baseline single-agent model.

4.5 Robustness and Generalization

Robustness. We evaluate Qwen3-8B distilled from Qwen3-32B using the GSM8K dataset on TruthfulQA [34]. It measures the model's ability to maintain factual accuracy and coherent reasoning. As shown in Table 1b, all distillation methods improve performance over the base model, indicating their robustness and resilience to catastrophic forgetting. In particular, PAD achieves the highest scores across all metrics, indicating more robust reasoning and better retention of factual correctness. These results suggest that MAS distillation not only enhances average accuracy but also increases robustness, allowing the student model to generalize more reliably to unseen or challenging tasks.

Open-ended Generalization. We evaluate whether reasoning skills learned from mathematical problem-solving transfer to other complex reasoning tasks such as open-ended questions. We train on a single source GSM8K and evaluate on three out-of-domain (OOD) datasets spanning diverse tasks: HotpotQA [65] (multi-hop reasoning), QASPER [10] (long-context understanding), and QMSum [74] (summarization).

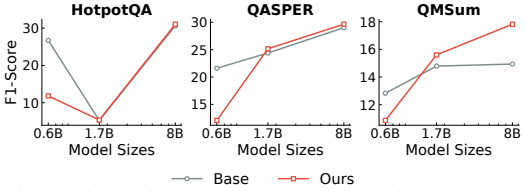

Figure 5: Performance on 3 open-ended datasets.

For comprehensive open-ended assessment, we leverage F1 scores and ROUGE-1/2/L [33] to measure lexical overlap between each prediction and the ground-truth, and BERTScore [69] for semantic similarity. Figure 5 shows that AgentArk significantly enhances cross-dataset reasoning transfer, particularly for larger models, which consistently improve performance on diverse OOD tasks such as multi-hop QA, long-context understanding, and summarization. This indicates that AgentArk strengthens general reasoning capabilities rather than merely fitting dataset-specific pattern. More details on other metrics are in Appendix D.1.

4.6 Extension to Multimodal LLMs, Heterogeneous Agents, and MAS Protocols

Distillation to Multimodal LLMs (MLLMs). We perform distillation from Qwen2.5-VL-32B-Instruct into the smaller 3B version [1]. Figure 6 shows that MAS distillation remains effective for MLLMs, despite being trained on text-only reasoning datasets. Across both benchmarks, PAD consistently achieves the strongest or near-

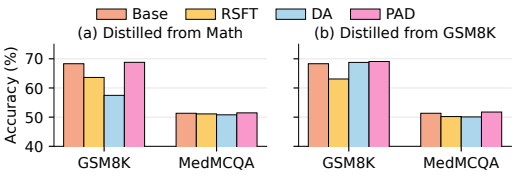

Figure 6: Multimodal distillation results.

Table 2: AgentArk in different MAS protocols and heterogeneous settings.

(a) Performance in different MAS protocols.

MAS Method	Single	RSFT	DA	PAD
AgentVerse	88.17 / 59.65	88.12 / 59.70 (-0.05 / +0.05)	88.48 / 59.85 (+0.31 / +0.20)	88.53 / 59.78 (+0.36 / +0.13)
Self-Consistency	88.17 / 59.65	88.20 / 59.58 (+0.03 / -0.07)	88.57 / 59.86 (+0.40 / +0.21)	88.58 / 59.81 (+0.41 / +0.16)

(b) Performance in heterogeneous settings.

Student Model	Single	RSFT	DA	PAD
Qwen3-1.7B	69.14	69.67	69.63	69.69
Gemma-7B	51.10	52.95	53.77	53.21

strongest performance, indicating that distilling process-level reasoning signals generalizes better than RSFT and DA. Notably, the absolute gains are modest compared to text-only settings, which is expected since the model is not explicitly trained on MLLM reasoning. This indicates that AgentArk captures model-agnostic reasoning patterns that are transferable beyond the training modality.

Generalization across Multi-Agent Protocols. AgentArk is designed to be algorithm-agnostic, as it distills reasoning dynamics from multi-agent interactions rather than relying on a specific coordination mechanism. To validate this property, we extend beyond the standard debate setting and evaluate two additional MAS paradigms: Self-Consistency [60] and AgentVerse [6]. We distill from Qwen3-32B to Qwen3-8B using GSM8K as the training dataset, and evaluate on both GSM8K and MedMCQA. As shown in Table 2a, AgentArk consistently improves student performance across all MAS variants and distillation strategies, showing that our approach captures transferable reasoning patterns at the level of interaction dynamics, rather than overfitting to a particular multi-agent algorithm.

Heterogeneous Settings. Finally, we study a more challenging *heterogeneous* setting, where agents are instantiated from *different* model families rather than a homogeneous backbone. Table 2b shows that AgentArk consistently improves performance over the single-agent baseline, demonstrating strong robustness to heterogeneous teacher compositions. confirming that AgentArk generalizes well across different multi-agent configurations. More details are in Appendix G.3.

5 Discussion

Our experiments reveal several insights and avenues for future research.

(1) MAS distillation strategies: Structured distillation, particularly PRM-guided methods, effectively transfers complex reasoning behaviors to smaller or multimodal models. Future work could explore adaptive distillation strategies that dynamically select reasoning trajectories based on task complexity, model capacity, or ensemble diversity. Additionally, methods that weigh intermediate self-corrections differently from final answers may further improve learning efficiency and reasoning fidelity.

(2) Process modeling and policy optimization: Larger PRMs yield stronger gains even for small target models. Future research could investigate modular or hierarchical PRMs, which allow selective guidance for different reasoning components, and alternative policy optimization techniques that balance stability and scalability beyond PPO and GRPO.

(3) Scaling agent and data quality: Increasing the number of debating agents benefits larger models but shows diminishing returns for smaller models. Similarly, training data volume alone is insufficient; high-quality reasoning trajectories are crucial. Future work could focus on adaptive ensemble scaling or selective trajectory sampling that optimizes the trade-off between diversity and learnability.

(4) Implication for foundation models: Our study suggests that it is promising to leverage MAS to augment single LLMs, which will dramatically reduce cost with improved performance. Moreover, the distillation on small models implies that small language models can be significantly enhanced by MAS, highlighting promise for lightweight and cost-effective deployment of foundation models.

Limitations. This work has the following limitations. First, experiments are limited to a subset of reasoning benchmarks and multimodal models, evaluating additional tasks and modalities could help assess broader applicability. Second, we focus on a specific set of distillation pipelines. Exploring alternative or hybrid approaches may provide further insights.

6 Conclusion

We introduced AgentArk, a framework for distilling multi-agent reasoning into a single agent. Our experiments demonstrate that structured distillation, particularly PRM-guided methods, enables smaller models to approximate complex reasoning behaviors while maintaining efficiency and generalization across diverse tasks. With the increasing attention to LLM efficiency and multi-agent systems, we hope our findings can provide insights for future research.

Acknowledgment

This work was supported by Amazon Research Award, Spring 2025. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of Amazon. We also highlight the computing supported by Modal Academic Compute Grant and William & Mary Research Computing for providing computational resources and/or technical support that have contributed to the results reported within this paper. URL: https://www.wm.edu/it/rc.

References

[1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

[2] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.

[3] Huaben Chen, Wenkang Ji, Lufeng Xu, and Shiyu Zhao. Multi-agent consensus seeking via large language models. *arXiv preprint arXiv:2310.20151*, 2023.

[4] Justin Chen, Swarnadeep Saha, and Mohit Bansal. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7066–7085, 2024.

[5] Justin Chih-Yao Chen, Swarnadeep Saha, Elias Stengel-Eskin, and Mohit Bansal. Magdi: Structured distillation of multi-agent interaction graphs improves reasoning in smaller language models. *arXiv preprint arXiv:2402.01620*, 2024.

[6] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, et al. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*, 2023.

[7] Weize Chen, Jiarui Yuan, Chen Qian, Cheng Yang, Zhiyuan Liu, and Maosong Sun. Optima: Optimizing effectiveness and efficiency for llm-based multi-agent system. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11534–11557, 2025.

[8] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

[9] Petru L Curşeu, Nicoleta Meslec, Helen Pluut, and Gerardus JM Lucas. Cognitive synergy in groups and group-to-individual transfer of decision-making competencies. *Frontiers in psychology*, 6:1375, 2015.

[10] Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A Smith, and Matt Gardner. A dataset of information-seeking questions and answers anchored in research papers. *arXiv preprint arXiv:2105.03011*, 2021.

[11] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*, 2023.

[12] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv:2407.21783*, 2024.

[13] Sugyeong Eo, Hyeonseok Moon, Evelyn Hayoon Zi, Chanjun Park, and Heuiseok Lim. Debate only when necessary: Adaptive multiagent collaboration for efficient llm reasoning. *arXiv preprint arXiv:2504.05047*, 2025.

[14] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*, 2024.

[15] Shanshan Han, Qifan Zhang, Yuhang Yao, Weizhao Jin, and Zhaozhuo Xu. Llm multi-agent systems: Challenges and open problems. *arXiv preprint arXiv:2402.03578*, 2024.

[16] Pengfei He, Yue Xing, Shen Dong, Juanhui Li, Zhenwei Dai, Xianfeng Tang, Hui Liu, Han Xu, Zhen Xiang, and Charu C Aggarwal. Comprehensive vulnerability analysis is necessary for trustworthy llm-mas. *arXiv preprint arXiv:2506.01245*, 2025.

[17] Mahmood Hegazy. Diversity of thought elicits stronger reasoning capabilities in multi-agent debate frameworks. *arXiv preprint arXiv:2410.12853*, 2024.

[18] Hyeon Hwang, Yewon Cho, Chanwoong Yoon, Yein Park, Minju Song, Kyungjae Lee, Gangwoo Kim, and Jaewoo Kang. Assessing llm reasoning steps via principal knowledge grounding. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 19925–19948, 2025.

[19] Ashrey Ignise and Yashika Vahi. Applications of multi-agent systems. *Authorea Preprints*, 2024.

[20] InternLM2. Internlm2 technical report, 2024.

[21] Minki Kang, Jongwon Jeong, Seanie Lee, Jaewoong Cho, and Sung Ju Hwang. Distilling llm agent into small models with retrieval and code tools. *arXiv preprint arXiv:2505.17612*, 2025.

[22] Zixuan Ke, Yifei Ming, Austin Xu, Ryan Chin, Xuan-Phi Nguyen, Prathyusha Jwalapuram, Semih Yavuz, Caiming Xiong, and Shafiq Joty. Mas-orchestra: Understanding and improving multi-agent reasoning through holistic orchestration and controlled benchmarks. *arXiv preprint arXiv:2601.14652*, 2026.

[23] Jiin Kim, Byeongjun Shin, Jinha Chung, and Minsoo Rhu. The cost of dynamic reasoning: Demystifying ai agents and test-time scaling from an ai infrastructure perspective. *arXiv preprint arXiv:2506.04301*, 2025.

[24] Yubin Kim, Ken Gu, Chanwoo Park, Chunjong Park, Samuel Schmidgall, A Ali Heydari, Yao Yan, Zhihan Zhang, Yuchen Zhuang, Mark Malhotra, et al. Towards a science of scaling agent systems. *arXiv preprint arXiv:2512.08296*, 2025.

[25] Tian Lan, Wenwei Zhang, Chengqi Lyu, Shuaibin Li, Chen Xu, Heyan Huang, Dahua Lin, Xian-Ling Mao, and Kai Chen. Training language models to critique with multi-agent feedback. *arXiv preprint arXiv:2410.15287*, 2024.

[26] Jinu Lee and Julia Hockenmaier. Evaluating step-by-step reasoning traces: A survey. *arXiv preprint arXiv:2502.12289*, 2025.

[27] Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. Camel: Communicative agents for" mind" exploration of large language model society. *Advances in neural information processing systems*, 36:51991–52008, 2023.

[28] Weizhen Li, Jianbo Lin, Zhuosong Jiang, Jingyi Cao, Xinpeng Liu, Jiayu Zhang, Zhenqiang Huang, Qianben Chen, Weichen Sun, Qiexiang Wang, et al. Chain-of-agents: End-to-end agent foundation models via multi-agent distillation and agentic rl. *arXiv preprint arXiv:2508.13167*, 2025.

[29] Xiaoxiao Li. When single-agent with skills replace multi-agent systems and when they fail. *arXiv preprint arXiv:2601.04748*, 2026.

[30] Xinyi Li, Sai Wang, Siqi Zeng, Yu Wu, and Yi Yang. A survey on llm-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinagearth*, 1(1):9, 2024.

[31] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 17889–17904, 2024.

[32] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.

[33] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.

[34] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 3214–3252, 2022.

[35] Yi-Cheng Lin, Kang-Chieh Chen, Zhe-Yan Li, Tzu-Heng Wu, Tzu-Hsuan Wu, Kuan-Yu Chen, Hung-yi Lee, and Yun-Nung Chen. Creativity in llm-based multi-agent systems: A survey. *arXiv preprint arXiv:2505.21116*, 2025.

[36] Jiaqi Liu, Chengkai Xu, Peng Hang, Jian Sun, Mingyu Ding, Wei Zhan, and Masayoshi Tomizuka. Language-driven policy distillation for cooperative driving in multi-agent reinforcement learning. *IEEE Robotics and Automation Letters*, 2025.

[37] Jun Liu, Zhenglun Kong, Peiyan Dong, Changdi Yang, Tianqi Li, Hao Tang, Geng Yuan, Wei Niu, Wenbin Zhang, Pu Zhao, et al. Structured agent distillation for large language model. *arXiv preprint arXiv:2505.13820*, 2025.

[38] Wendy Yaqiao Liu, Rui Jerry Huang, Anastasia Miin, and Lei Ding. Adaptive coopetition: Leveraging coarse verifier signals for resilient multi-agent llm reasoning. In *The 14th International Joint Conference on Natural Language Processing and The 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 145–155, 2025.

[39] Tie Ma, Yixi Chen, Vaastav Anand, Alessandro Cornacchia, Amândio R Faustino, Guanheng Liu, Shan Zhang, Hongbin Luo, Suhaib A Fahmy, Zafar A Qazi, et al. Maestro: Multi-agent evaluation suite for testing, reliability, and observability. *arXiv preprint arXiv:2601.00481*, 2026.

[40] Sagnik Mukherjee, Abhinav Chinta, Takyoung Kim, Tarun Anoop Sharma, and Dilek Hakkani-Tür. Premise-augmented reasoning chains improve error identification in math reasoning with llms. *arXiv preprint arXiv:2502.02362*, 2025.

[41] Joaquin Navajas, Tamara Niella, Gerry Garbulsky, Bahador Bahrami, and Mariano Sigman. Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour*, 2(2):126–132, 2018.

[42] Thi-Nhung Nguyen, Linhao Luo, Thuy-Trang Vu, and Dinh Phung. The social cost of intelligence: Emergence, propagation, and amplification of stereotypical bias in multi-agent systems. *arXiv preprint arXiv:2510.10943*, 2025.

[43] Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pages 248–260. PMLR, 2022.

[44] Jiqun Pan, Zhenke Duan, Jiani Tu, Anzhi Cheng, and Yanqing Wang. Knowledge graph-guided multi-agent distillation for reliable industrial question answering with datasets. *arXiv preprint arXiv:2510.06240*, 2025.

[45] Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Xuanhe Zhou, Yufei Huang, Chaojun Xiao, et al. Tool learning with foundation models. *ACM Computing Surveys*, 57(4):1–40, 2024.

[46] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.

[47] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

[48] Xi Shi, Mengxin Zheng, and Qian Lou. Learning latency-aware orchestration for parallel multi-agent systems. *arXiv preprint arXiv:2601.10560*, 2026.

[49] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv:2503.19786*, 2025.

[50] Wataru Toyokawa, Andrew Whalen, and Kevin N Laland. Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 3(2):183–193, 2019.

[51] David Wan, Justin Chen, Elias Stengel-Eskin, and Mohit Bansal. Mamm-refine: A recipe for improving faithfulness in generation with multi-agent collaboration. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9882–9901, 2025.

[52] Guancheng Wan, Lucheng Fu, Haoxin Liu, Yiqiao Jin, Hui Yi Leong, Eric Hanchen Jiang, Hejia Geng, Jinhe Bi, Yunpu Ma, Xiangru Tang, et al. Beyond magic words: Sharpness-aware prompt evolving for robust large language models with tare. In *ICLR*, 2026.

[53] Jinlong Wang, Yunting Wu, Xiaoyun Xiong, Yuanyuan Zhang, Zhihan Lyu, Ahmed Ghoneim, and Haoran Zhao. Fedlma: A federated learning framework integrating llm-based multi-agent reasoning with knowledge distillation. *IEEE Transactions on Consumer Electronics*, 2025.

[54] Kun Wang, Guibin Zhang, ManKit Ye, Xinyu Deng, Dongxia Wang, Xiaobin Hu, Jinyang Guo, Yang Liu, and Yufei Guo. Mas^2: Self-generative, self-configuring, self-rectifying multi-agent systems. *arXiv preprint arXiv:2509.24323*, 2025.

[55] Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9426–9439, 2024.

[56] Peng-Yuan Wang, Tian-Shuo Liu, Chenyang Wang, Ziniu Li, Yidi Wang, Shu Yan, Chengxing Jia, Xu-Hui Liu, Xinwei Chen, Jiacheng Xu, et al. A survey on large language models for mathematical reasoning. *ACM Computing Surveys*, 2025.

[57] Song Wang, Zhen Tan, Zihan Chen, Shuang Zhou, Tianlong Chen, and Jundong Li. Anymac: Cascading flexible multi-agent collaboration via next-agent prediction. *arXiv preprint arXiv:2506.17784*, 2025.

[58] Xiao Wang, Jia Wang, Yijie Wang, Pengtao Dang, Sha Cao, and Chi Zhang. Mars: toward more efficient multi-agent collaboration for llm reasoning. *arXiv preprint arXiv:2509.20502*, 2025.

[59] Xinchen Wang, Pengfei Gao, Xiangxin Meng, Chao Peng, Ruida Hu, Yun Lin, and Cuiyun Gao. Aegis: An agent-based framework for general bug reproduction from issue descriptions. *arXiv preprint arXiv:2411.18015*, 2024.

[60] 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.

[61] Yiyang Wang, Chen Chen, Tica Lin, Vishnu Raj, Josh Kimball, Alex Cabral, and Josiah Hester. Companioncast: A multi-agent conversational ai framework with spatial audio for social co-viewing experiences. *arXiv:2512.10918*, 2025.

[62] Yiyang Wang, Yiqiao Jin, Alex Cabral, and Josiah Hester. Mascot: Towards multi-agent socio-collaborative companion systems. *arXiv:2601.14230*, 2026.

[63] Tianxin Wei, Ting-Wei Li, Zhining Liu, Xuying Ning, Ze Yang, Jiaru Zou, Zhichen Zeng, Ruizhong Qiu, Xiao Lin, Dongqi Fu, et al. Agentic reasoning for large language models. *arXiv preprint arXiv:2601.12538*, 2026.

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

[65] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380, 2018.

[66] Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

[67] Peiying Yu, Guoxin Chen, and Jingjing Wang. Table-critic: A multi-agent framework for collaborative criticism and refinement in table reasoning. *arXiv preprint arXiv:2502.11799*, 2025.

[68] Yurun Yuan and Tengyang Xie. Reinforce LLM reasoning through multi-agent reflection. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 73701–73731. PMLR, 13–19 Jul 2025.

[69] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2019.

[70] Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems*, 43(6):1–47, 2025.

[71] Siyan Zhao, Zhihui Xie, Mengchen Liu, Jing Huang, Guan Pang, Feiyu Chen, and Aditya Grover. Self-distilled reasoner: On-policy self-distillation for large language models. *arXiv preprint arXiv:2601.18734*, 2026.

[72] Zhonghan Zhao, Ke Ma, Wenhao Chai, Xuan Wang, Kewei Chen, Dongxu Guo, Yanting Zhang, Hongwei Wang, and Gaoang Wang. Do we really need a complex agent system? distill embodied agent into a single model. *arXiv preprint arXiv:2404.04619*, 2024.

[73] Xinyi Zheng, Ningke Li, Xiaokun Luan, Kailong Wang, Ling Shi, Meng Sun, and Haoyu Wang. Beyond correctness: Exposing llm-generated logical flaws in reasoning via multi-step automated theorem proving. *arXiv preprint arXiv:2512.23511*, 2025.

[74] Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, et al. Qmsum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921, 2021.

[75] Yiwu Zhong, Zi-Yuan Hu, Yin Li, and Liwei Wang. Rethinking chain-of-thought reasoning for videos. *arXiv:2512.09616*, 2025.

[76] Xiaofeng Zhou, He-Yan Huang, and Lizi Liao. Debate, reflect, and distill: Multi-agent feedback with tree-structured preference optimization for efficient language model enhancement. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 9122–9137, 2025.

Appendix

Table of Contents

A Data Generation Methods and Pipeline
- A.1 Dataset Overview
- A.2 Multi-Agent Data Generation
- A.3 Correctness Filtering
- A.4 Data Selection for Distillation with Data Augmentation

B Methods
- B.1 Process Reward Model: Contrastive Loss Design
- B.2 Ablation: PPO Comparison

C Ablation Studies on PAD
- C.1 PRM Modeling and Policy Separation
- C.2 PPO vs. GRPO

D Generalization
- D.1 Generalization on Out-of-Domain Datasets
- D.2 Data Scaling
- D.3 Computation Cost

E Case Study
- E.1 Perplexity Calculation

F Combined Distillation Strategies with Two-Step Training
- F.1 Two-Step Training Protocol

G Distillation Results
- G.1 Supervised Fine-tuning
- G.2 Reasoning Based Distillation Results

A Data Generation Methods and Pipeline

This section provides detailed descriptions of the data generation pipeline, dataset composition, and statistics used in our experiments.

A.1 Dataset Overview

We construct the multi-agent distillation dataset using four core benchmarks, focusing on training the student models:

- **Mathematical reasoning:** GSM8K and MATH, covering arithmatic, algebraic, and multi-step symbolic reasoning.
- **Augmented math reasoning:** MetaMathQA, providing additional multi-step problems and varied solution strategies.
- **Domain-specific knowledge:** MedMCQA, focusing on medical exam-style multiple-choice questions.

The remaining benchmarks, HotpotQA, QAPER, and QMSum, are reserved exclusively for zero-shot generalization evaluation and are not used during training or data augmentation.

For the training datasets, after multi-agent data generation, correctness filtering, and diversity-based selection, the final distillation dataset contains approximately 342k unique input questions and 2M reasoning trajectories.

Table 3 reports detailed statistics for each training dataset, including the number of questions, retained debates, and augmented trajectories.

Table 3: Dataset statistics for multi-agent distillation. Q indicates the number of unique training questions, while T denotes the number of correct reasoning trajectories extracted per dataset.

Source Model	GSM8K(Q / T)	Math(Q / T)	MetaMathQA(Q / T)	Medmcqa(Q / T)
Qwen3-8B	7473/44,838	500/3000	151k/906k	183k/1.1M
Qwen3-32B	7473/44,838	500/3000	151k/906k	183k/1.1M
Gemma3-27B-It	7473/44,838	500/3000	151k/906k	183k/1.1M

A.2 Multi-Agent Data Generation

To generate rich reasoning supervision, we adopt a multi-agent debate mechanism to produce diverse solution trajectories. For each input problem x, we initialize a set of $n = 5$ agents, plus a agent that summarize the final answer at the end. $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ that share the same underlying teacher language model but operate with independent generation contexts. Agents engage in a debate for up to $K = 3$ rounds. In the first round, each agent independently produces an initial reasoning trace and final answer. In subsequent rounds, each agent observes the reasoning traces generated by other agents in the previous round and is encouraged to revise its own solution by identifying potential errors, alternative solution paths, or overlooked details. Formally, in round k, agent a_i generates a reasoning trajectory $\tau_{i,k}$ conditioned on the problem x and the peer trajectories $\{\tau_{j,k-1}\}_{j \neq i}$. The complete debate log for an input x is denoted as

$$\mathcal{L}_x = \{\tau_{i,k} \mid i \in [1, n], k \in [1, K]\}.$$

This interaction protocol promotes reasoning behaviors such as self-correction, hypothesis revision, and cross-verification, yielding a diverse set of candidate solutions that go beyond single-pass generation.

A.3 Correctness Filtering

For each input problem x, we collect responses generated by a set of agents $\mathcal{A}$, where each agent produces a final answer along with its associated reasoning trajectory. We first apply a *correctness filtering* step to identify agents whose final answers are valid.

To ensure consistent and scalable evaluation across all tasks, we use Qwen2.5-72B-Instruct as an automatic verifier to assess whether a candidate final answer matches the ground-truth solution under task-specific evaluation rules (e.g., exact match or normalized equivalence for GSM8K-style problems). This model is prompted to perform answer verification only and does not take the reasoning trajectory into account during this stage.

An agent $a \in \mathcal{A}$ is included in the set of successful contributors $\mathcal{A}_{\text{correct}}$ if its final answer is verified as correct by the verifier model. This step ensures that all candidate responses used for distillation are answer-correct, independent of their reasoning styles.

Formally,

$$\mathcal{A}_{\text{correct}} = \{a \in \mathcal{A} \mid \text{Answer}(a, x) \text{ is verified as correct}\}. \tag{5}$$

If $|\mathcal{A}_{\text{correct}}| < 2$, the corresponding problem instance is excluded from the data augmentation process, as it does not provide sufficient diversity among correct solutions.

A.4 Data Selection for Distillation with Data Augmentation

A.4.1 Divergent-Reasoning Selection via LLM-based Judgment

Although all agents in $\mathcal{A}_{\text{correct}}$ arrive at the same correct answer, their reasoning processes may differ substantially. To capture this variation, we adopt a Correctness-First Diverse Extraction strategy that emphasizes reasoning divergence under answer agreement.

For each problem x, we consider the set of reasoning traces produced by agents in $\mathcal{A}_{\text{correct}}$. We then employ Qwen2.5-72B-Instruct as an auxiliary judge to identify responses whose reasoning patterns are meaningfully distinct while preserving correctness.

The judge model is provided with the problem x, the ground-truth answer, and multiple candidate reasoning traces from $\mathcal{A}_{\text{correct}}$. It is instructed to select responses that satisfy the following criteria:

1. The final answer is correct.
2. The reasoning process exhibits structural differences compared to other correct responses (e.g., different decomposition orders, intermediate representations, or solution paths).

Notably, the judge model is not used to rank responses by quality or correctness. Instead, it operates solely to assess reasoning diversity conditioned on answer correctness, which reduces bias toward specific reasoning styles.

A.4.2 Final Augmented Dataset Construction

The selected responses are incorporated into the distillation dataset as augmented supervision signals. By construction, the resulting dataset satisfies two properties:

- **Answer consistency:** all augmented samples preserve the correct final answer.
- **Reasoning diversity:** multiple valid reasoning trajectories are retained for the same input.

This data selection procedure enables the student model to learn from a richer set of correct problem-solving behaviors without introducing incorrect or contradictory supervision.

B Methods

B.1 Process Reward Model: Contrastive Loss Design

Denoting $\sigma(\cdot)$ as the $\mathrm{sigmoid}$ function, we design the Process Reward Model (PRM) loss as a **contrastive loss** rather than standard binary cross-entropy. This choice encourages the model to assign higher rewards to reasoning steps that are more consistent with the multi-agent debate consensus, reflecting relative correctness rather than absolute labels.

Specifically, given a positive reasoning step r_t^+ and a set of negative steps $\{r_t^-\}$ sampled from other agents' reasoning trajectories, the PRM loss is defined as:

$$\mathcal{L}_{\text{PRM}}(\phi) = -\sum_t \log \mathrm{softmax}\left(\frac{\sigma(R_\phi(r_t^+)), \{\sigma(R_\phi(r^-))\}_{r^- \in \mathcal{N}_t}}{\tau}\right), \tag{6}$$

where τ is a temperature hyperparameter controlling the sharpness of the contrastive distribution. This formulation encourages the PRM to score more consistent reasoning steps higher while down-weighting less consistent or contradictory steps.

B.2 Ablation: PPO Comparison

In addition to GRPO, we also experimented with standard Proximal Policy Optimization (PPO) [46] to verify the benefits of group-wise relative updates. The overall setup is identical to GRPO [47] training, including the use of the Process Reward Model (PRM) to provide step-level supervision. The student policy π_θ is fine-tuned using the conventional PPO clipped objective:

$$\mathcal{L}_{\text{PPO}}(\theta) = \mathbb{E}_{x\sim\mathcal{D}, o\sim\pi_{\text{old}}}\Big[\min\left(\rho(\theta)\hat{A}, \text{clip}(\rho(\theta), 1-\epsilon, 1+\epsilon)\hat{A}\right) - \beta\mathbb{D}_{\text{KL}}(\pi_\theta\|\pi_{\text{ref}})\Big], \tag{7}$$

where $\rho(\theta) = \frac{\pi_\theta(o|x)}{\pi_{\text{old}}(o|x)}$ is the probability ratio and $\hat{A}$ is the advantage derived from PRM scores:

$$\hat{A} = \frac{R_\phi(o) - \mu_R}{\sigma_R}, \tag{8}$$

with μ_R and σ_R denoting the mean and standard deviation of PRM scores across sampled outputs for a given input x.

Unlike GRPO, PPO treats each sample independently rather than comparing outputs within a group, which can lead to slower convergence and reduced reward consistency in multi-step reasoning tasks.

The comparison between GRPO and PPO in our ablation (Table X) demonstrates that GRPO achieves more stable learning and higher final reasoning accuracy, validating the benefit of group-relative updates.

C Ablation Studies on PAD

C.1 PRM Modeling and Policy Separation

We study the role of PRM by explicitly separating the model used to learn process rewards from the model used to train the final policy. In our setup, PRM is first trained independently, and its learned process-level guidance is then used to supervise the distillation of a target policy model. All experiments distill from the same Qwen3-32B, while the PRM and the policy are instantiated with different parameter scales to analyze the effect of PRM capacity under this decoupled training scheme. As shown in Table 4, using a smaller PRM model consistently leads to limited improvement, even when paired with the same policy model. In contrast, assigning a larger one yields gains across the GSM8K and MedMCQA test sets, even when the target policy remains lightweight. These results indicate that effective distillation relies more critically on the modeling size of the PRM than on the scale of the policy itself, highlighting the importance of explicit process modeling and role separation.

Table 4: Ablation on PRM and LLM role separation. All models are distilled from Qwen3-32B.

Test set	PRM	LLM	GSM8K	MATH	MedMCQA	MMQA	Avg
GSM8K	0.6B	0.6B	42.84	42.68	42.46	42.53	42.63
	8B	0.6B	42.84	42.53	42.91	42.23	42.63
	0.6B	8B	88.48	88.55	88.40	88.63	88.52
	8B	8B	88.63	88.70	88.48	88.56	88.59
MedMCQA	0.6B	0.6B	32.39	32.58	32.36	32.54	32.47
	8B	0.6B	32.49	32.56	32.35	32.20	32.40
	0.6B	8B	59.50	59.74	59.74	59.65	59.66
	8B	8B	59.77	59.81	59.86	59.77	59.80

C.2 PPO vs. GRPO

To optimize the target policy during PRM-guided distillation, we compare Proximal Policy Optimization (PPO) [46] (details in Section B.2) with GRPO. PRMs in both settings are trained identically in a first stage and then fixed. The two methods differ only in the second stage, where the target language model is fine-tuned using either PPO or GRPO under the same PRM-based reward signals. As shown in Table 5, PPO achieves marginally higher accuracy in most settings. This advantage is expected, as PPO employs a learned value function to provide a lower-variance, state-dependent baseline for policy updates, which can lead to more stable optimization under imperfect PRM rewards. In contrast, GRPO removes the value function and relies on group-relative comparisons, resulting in slightly noisier updates but significantly reduced computational overhead. Despite this difference, GRPO achieves performance comparable to PPO across all benchmarks, making it a scalable and effective alternative for large-scale PRM-guided distillation.

Table 5: Ablation on PPO vs GRPO.

	Test set	GSM8K	MedMCQA	MMQA	Math	Avg
PPO	GSM8K	53.37	53.15	53.24	52.62	53.10
	MedMCQA	49.56	50.08	49.77	49.41	49.71
GRPO	GSM8K	52.71	52.48	52.62	52.64	52.61
	MedMCQA	49.77	50.02	49.26	49.33	49.60

D Generalization

D.1 Generalization on Out-of-Domain Datasets

To evaluate the robustness of AgentArk, we conduct experiments across three out-of-domain (OOD) datasets spanning diverse domains: HotpotQA [65] (multi-hop reasoning), QASPER [10] (long-context understanding), and QMSum [74] (summarization). The results in Figure 5 and Appendix Figure 7 & 8 reveal a clear scaling law for reasoning transferability.

The most substantial gains are observed in the larger 8B parameter class. Qwen3-8B and Llama3-8B demonstrate superior generalizability, effectively internalizing multi-agent thought processes to

achieve consistent performance lifts across all OOD tasks. This 'reasoning lift' is most pronounced in the complex QMSum task, where Qwen3-8B improves its F1-Score from 14.94 to 17.82 and its ROUGE-L Score from 15.72 to 17.41. Similarly, Llama3-8B sees its F1-Score increase from 13.05 to 14.92 on QMSum. These gains suggest that AgentArk is particularly potent for complex tasks such as high-level synthesis and human-centered dialog understanding, where the model must perform multi-step reasoning and aggregate information from extensive documents multiple perspectives.

Even mid-sized models like Qwen3-1.7B begin to show this positive trajectory, with F1 scores rising from 24.38 to 25.16 on QASPER and from 14.79 to 15.60 on QMSum. This indicates that once a model surpasses a critical capacity threshold, AgentArk serves as an enhancer for its inherent reasoning capabilities. In contrast, ultra-small models like Qwen3-0.6B have limited capacity to balance new reasoning patterns with its existing pre-trained knowledge base. Thus, AgentArk's multi-agent knowledge is most effective for models with sufficient cognitive capacity.

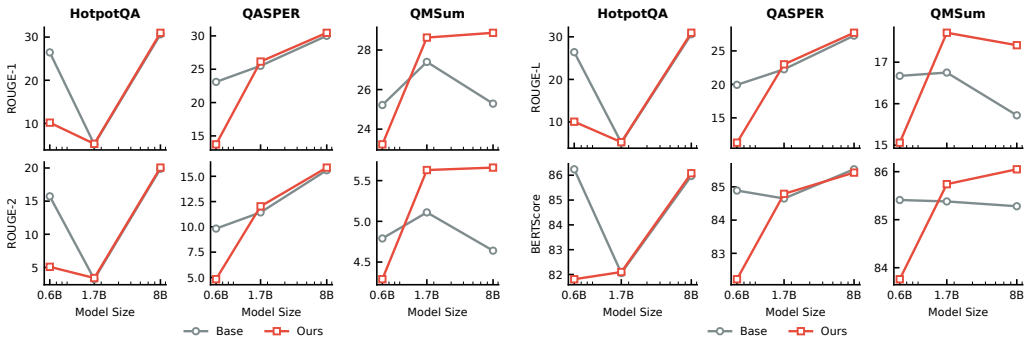

Figure 7: Scaling behavior of AgentArk vs. base models. Performance comparison on three out-of-distribution (OOD) datasets across the Qwen model family demonstrate the impact of increasing model size (0.6B to 8B) on generalizability.

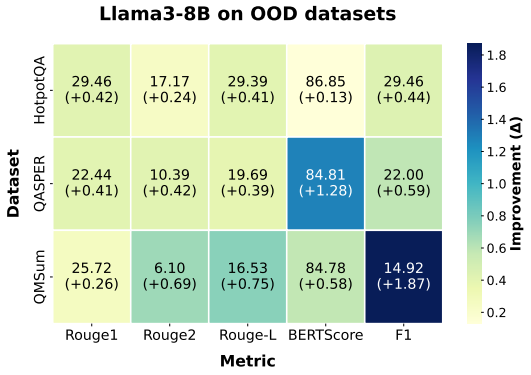

Figure 8: Performance of AgentArk and the base models on out-of-domain datasets

D.2 Data Scaling

We provide additional in Figure 9 data-scaling results using the MedMCQA dataset, complementing the scaling analysis presented in Section 4.3 of the main paper.

D.3 Computation Cost

While AgentArk significantly reduces inference-time cost by eliminating multi-agent coordination, its training pipeline, particularly Process-Aware Distillation (PAD) introduces additional computational overhead. We report concrete training costs for different distillation strategies to provide a realistic assessment of their practicality.

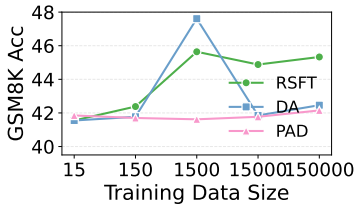

Figure 9: Data scaling behavior on MedMCQA for distillation from Qwen3-32B to Qwen3-0.6B, showing student model performance as a function of the amount of training data.

Experimental Setup. All experiments are conducted on NVIDIA H100 GPUs (80GB). Unless otherwise noted, we are reporting the student model has 8B parameters and is trained with a global batch size of 4 and identical sequence length across methods. All distillation approaches share the same multi-agent data generation cost, differences arise solely from the supervision signal and optimization procedure applied during student training.

Training Cost. Table 6 summarizes the additional training cost incurred by different distillation strategies.

Table 6: Training cost for different distillation strategies (8B student model).

Method	Additional Training Components	GPUs	Time
RSFT	Supervised fine-tuning on single reasoning traces	1 × H100	∼6 hours
Reasoning DA	Supervised fine-tuning on augmented multi-trajectory reasoning data	1 × H100	∼8 hours
PAD (PRM + GRPO)	PRM training + GRPO-based policy optimization	8 × H100	∼20 hours
PRM training	Step-level process reward modeling	8 × H100	∼8 hours
GRPO	Policy optimization with PRM reward	8 × H100	∼12 hours

Inference-Time Cost. In contrast to the increased offline training cost, inference with AgentArk requires only a single autoregressive generation of the distilled student model. Compared to multi-agent debate, which typically involves multiple agent invocations and iterative coordination, AgentArk significantly reduces inference latency and GPU usage, making it suitable for real-time and resource-constrained deployment settings.

Overall, AgentArk trades increased offline computation for substantial reductions in inference-time cost and deployment complexity, consistent with common practices in large-scale model distillation and reinforcement learning.

E Case Study

Single Agent Reasoning

Let's think step by step... Step 1: Green eggs = 5. Step 2: Blue eggs = 2 * 5 = 10. Step 3: Pink eggs = 10 - 1 = 9. Step 4: Yellow eggs = 1/3 * 9 = 3. Total = 5 + 10 + 9 + 3 = 27.
Wait, but the answer is 24... let me check again...
Recalculating yellow eggs: 9 / 3 = 3. Total still 27. Something is wrong...
Wait, but the answer is 24... let me check again...
Recalculating yellow eggs: 9 / 3 = 3. Total still 27. Something is wrong...
Wait, but the answer is 24... let me check again...
Recalculating yellow eggs: 9 / 3 = 3. Total still 27. Something is wrong...
...

Multi-Agent Reasoning

Agent 1: Step 1: Green eggs = 5 Step 2: Blue eggs = 2 * 5 = 10 Step 3: Pink eggs = 10 - 1 = 9 Step 4: Yellow eggs = 9 / 3 = 3 Total = 5 + 10 + 9 + 3 = 27
Agent 2: Step 1: Green eggs = 5 Step 2: Blue eggs = 10 Step 3: Pink eggs = 9 Step 4: Yellow eggs = 3 Total = 27 Verifies reasoning and intermediate calculations
Agent 3: Step 1: Green = 5 Step 2: Blue = 10 Step 3: Pink = 9 Step 4: Yellow = 3 Total = 27 Confirms each step independently...
Agent 4: Step 1: Green = 5 Step 2: Blue = 10 Step 3: Pink = 9 Step 4: Yellow = 3 Total = 27 Checks reasoning consistency
Agent 5: Step 1: Green = 5 Step 2: Blue = 10 Step 3: Pink = 9 Step 4: Yellow = 3 Total = 27 All calculations verified
Final Answer: 27

Reasoning Comparison

AgentArk:
The problem states that Cindy had 5 green eggs.
Step 1: She had twice as many blue eggs as green eggs, so she had 2 * 5 = 10 blue eggs.
Step 2: She had one fewer pink eggs than blue eggs, so she had 10 - 1 = 9 pink eggs.
Step 3: She had one-third as many yellow eggs as pink eggs, so she had 9 / 3 = 3 yellow eggs.
Step 4: The total number of eggs is 5 + 10 + 9 + 3 = 27 eggs.
Final Answer: 27
Reasoning: The problem states that Cindy had 5 green eggs. She had twice as many blue eggs as green eggs, so she had 2 * 5 = 10 blue eggs. She had one fewer pink eggs than blue eggs, so she had 10 - 1 = 9 pink eggs. She had one-third as many yellow eggs as pink eggs, so she had 9 / 3 = 3 yellow eggs. The total number of eggs is 5 + 10 + 9 + 3 = 27 eggs. The reasoning provided in the final answer is correct and matches the problem statement.

E.1 Perplexity Calculation

Definition. Given a reasoning sequence $x = (x_1, \ldots, x_T)$, the perplexity (PPL) of a language model parameterized by θ is defined as:

$$\mathrm{PPL}(x) = \exp\left(-\frac{1}{T}\sum_{t=1}^{T}\log p_\theta(x_t \mid x_{<t})\right), \tag{9}$$

where $p_\theta(x_t \mid x_{<t})$ denotes the conditional probability assigned by the model to the next token x_t given all previous tokens.

Reasoning-Token Perplexity. Unlike standard perplexity evaluation over full responses, we restrict the computation to *reasoning tokens* only. Specifically, we identify reasoning spans corresponding to intermediate reasoning steps (e.g., chain-of-thought segments) and exclude prompt tokens, question descriptions, and final answer tokens. Let $\mathcal{R} \subseteq \{1, \ldots, T\}$ denote the index set of reasoning tokens. The reasoning perplexity is then computed as:

$$\mathrm{PPL}_{\mathrm{reason}}(x) = \exp\left(-\frac{1}{|\mathcal{R}|}\sum_{t\in\mathcal{R}}\log p_\theta(x_t \mid x_{<t})\right). \tag{10}$$

Evaluation Protocol. We evaluate perplexity on held-out GSM8K samples that are not used during distillation. For each example, the gold reasoning trajectory is tokenized using the same tokenizer as the evaluated model. The model is run in teacher-forcing mode to obtain token-level log-likelihoods. Perplexity is computed per sample and then averaged across all evaluation samples.

Model Setup. We compute perplexity for distilled student models and the single-agent baseline using the same teacher forcing procedure. The Qwen3-32B model serves as the teacher, while Qwen3-0.6B is used as the student architecture. All models are evaluated with identical prompts and decoding-free forward passes to ensure fair comparison.

Interpretation. Lower reasoning perplexity indicates that the model assigns higher likelihood to coherent and structured reasoning steps, reflecting improved internal consistency and predictability of the reasoning process. This aligns with our objective of distilling multi-agent reasoning dynamics into a single model.

F Combined Distillation Strategies with Two-Step Training

In addition to evaluating individual distillation strategies, including reasoning-enhanced supervised fine-tuning (RSFT), reasoning data augmentation (DA), and process-aware distillation (PAD), we further investigate whether these methods can be effectively combined in a sequential manner.

F.1 Two-Step Training Protocol

We adopt a two-step training scheme that stacks DA on top of an existing distillation method. Specifically, we consider **RSFT+DA** and **PAD+DA**, where: (i) in the first stage, the student model is trained using RSFT or PAD following the same setup as the main experiments; and (ii) in the second stage, the resulting model is further fine-tuned with reasoning data augmentation to enhance reasoning diversity.

All experiments distill from a **Qwen3-32B** source model to a **Qwen3-1.7B** student model. In the second stage, we separately use reasoning data generated by the 32B model on **GSM8K** and **MATH**, denoted as *32B_gsm8k* and *32B_math*. The *base* setting corresponds to using DA data aligned with the same task as the first-stage training.

Table 7 summarizes the results of mixed distillation strategies on GSM8K and MedMCQA.

Table 7: Performance of combined distillation strategies with two-step training.

Method	Train / Test	**Base**	**gsm8k**	**math**
RSFT+DA	GSM8K	68.61	70.27	69.89
	MedMCQA	42.67	43.92	44.30
PAD+DA	GSM8K	70.03	70.69	70.14
	MedMCQA	42.72	43.68	43.49

Overall, stacking DA on top of RSFT or PAD leads to consistent but modest gains across benchmarks. While the improvements are incremental, the results suggest that our methods are mutually compatible, as reasoning data augmentation can be applied on top of different distillation strategies without degrading performance.

G Distillation Results

G.1 Supervised Fine-tuning

We evaluate the performance of distillation through standard SFT trained solely with ground-truth answer supervision, following common practice. Specifically, models are fine-tuned using only input–output pairs from the ground-truth data, without any additional reasoning annotations, auxiliary losses, or trajectory-level supervision. After training, we assess the performance both in-distribution performance and out-of-distribution generation under distribution shifts that require changes in reasoning.

The results is shown in Table 8. A consistent pattern emerges across model families and scales. While SFT occasionally yields moderate gains on MedMCQA, it fails to produce consistent or reliable improvements on GSM8K, and in many cases leads to performance degradation. Specifically, MedMCQA exhibits small but repeatable improvements when models are fine-tuned on medically oriented or structurally similar datasets (e.g., +6.8 for Gemma-7B, +3.4 for Qwen3-8B). This suggests that answer-only supervision can be beneficial when the target task shares surface-level structure or domain overlap with the training data, allowing models to exploit task-specific correlations. In contrast, GSM8K performance remains largely unimproved or even deteriorates under SFT across nearly all settings. Even when trained directly on GSM8K, models frequently underperform their base

Table 8: SFT performance

Test \ Train		base	gsm8k	medmcqa	metamathqa	math
gsm8k	llama3-8b	75.28	63.61	73.31	65.78	75.36
medmcqa		56.9	60.79	60.44	60.46	60.34
gsm8k	qwen3-8b	88.17	81.35	88.17	87.41	87.87
medmcqa		59.65	60.05	63.02	59.98	59.67
gsm8k	gemma-7b	51.1	56.56	48.52	36.39	50.04
medmcqa		49.29	48.98	57.06	48.15	48.43
gsm8k	qwen3-1.7b	69.14	59.59	70.66	66.34	67.32
medmcqa		42.34	43.72	49.53	43.68	42.46
gsm8k	qwen3-0.6b	41.93	37.91	44.35	21.38	31.31
medmcqa		32.2	32.9	38.9	34.09	32.27

counterparts. This behavior indicates that SFT struggles to induce transferable reasoning strategies required for multi-step mathematical problem solving, and instead encourages overfitting to shallow input–output mappings that do not generalize beyond the supervised distribution. Overall, these results highlight a key limitation of standard SFT: while answer-only supervision may yield localized gains on domain-specific benchmarks, it fails to support robust cross-task or reasoning-intensive generalization. This empirical evidence motivates the need for process-level supervision that exposes models to intermediate reasoning dynamics rather than only final outcomes.

G.2 Reasoning Based Distillation Results

Results for our three proposed methods are listed in Tables 9 to 13.

G.3 AgentArk in Heterogeneous Setting

We construct a MAS consisting of five heterogeneous agents: Qwen3-32B, Gemma-3-27B-IT, Devstral-Small-2-24B-Instruct-2512, QwQ-32B, and Phi-4-reasoning. We evaluate two student models, Qwen3-1.7B and Gemma-7B, under this setting. The results are summarized in Table 2b. Across both students and all distillation strategies, AgentArk consistently improves performance over the single-agent baseline, demonstrating strong robustness to heterogeneous teacher compositions.

In particular, for Qwen3-1.7B, RSFT, DA, and PAD improve performance by +0.53, +0.49, and +0.55 points respectively. For Gemma-7B, the gains are even more substantial, reaching +1.85, +2.67, and +2.11 points respectively. These results indicate that heterogeneous teacher ensembles can provide complementary reasoning signals that further benefit distillation, especially for larger-capacity student models that can better absorb diverse supervision.

Overall, the performance trends are consistent with those observed in homogeneous MAS settings, confirming that AgentArk generalizes well across different multi-agent configurations. Importantly, the improvements remain stable across architectures and scales, suggesting that our approach effectively leverages diverse reasoning behaviors rather than relying on homogeneous consensus among agents.

H AI Assistants Usage

AI assistants were used as auxiliary tools in preparing this manuscript, primarily for language refinement, clarity and organization. The experimental design and methodological choices were made by the authors.

I Broader Impact

AgentArk offers meaningful societal benefits by enabling the advantages of multi agent reasoning through efficient distillation, rather than relying on expensive test time orchestration. This significantly reduces latency and deployment costs for reasoning intensive applications where traditional multi

Table 9: Distillation results on Qwen3-8B. Rows denote training datasets and columns denote test benchmarks (see diagonal header). Results compare RSFT, DA, and PAD across different source models. The original Qwen3-8B scores 88.17 on GSM8K and 59.65 on MedMCQA.

source model	test \ train	gsm8k	medmcqa	metamathqa	math
		RSFT			
qwen3-8b	gsm8k	87.11	86.88	87.72	87.19
	medmcqa	58.81	58.88	59.6	59.77
		DA			
	gsm8k	86.43	86.35	88.1	86.41
	medmcqa	58.33	58.83	57.49	58.35
		PAD			
	gsm8k	88.42	88.36	88.48	88.41
	medmcqa	59.71	59.81	59.71	59.73
		RSFT			
qwen3-32b	gsm8k	86.73	87.04	85.6	85.97
	medmcqa	60.03	59.74	57.95	59.96
		DA			
	gsm8k	87.11	87.49	89.57	86.2
	medmcqa	59.48	58.69	58.33	59.86
		PAD			
	gsm8k	89.05	89.02	89.15	88.7
	medmcqa	60.04	63.12	61.53	61.21
		RSFT			
gemma3-27b-it	gsm8k	86.96	86.35	87.72	87.64
	medmcqa	59.53	58.43	59.77	59.74
		DA			
	gsm8k	87.79	86.28	86.2	86.05
	medmcqa	59.43	58.5	57.18	59.72
		PAD			
	gsm8k	88.48	88.48	88.48	88.4
	medmcqa	59.96	59.84	59.86	59.79

agent inference is impractical, such as on device or resource constrained environments. By lowering the computational barrier, AgentArk broadens access to advanced agentic reasoning capabilities, potentially democratizing their use across a wider range of users and applications.

However, these benefits are accompanied by potential risks. Distilled student models may inherit undesirable behaviors from teacher models, including biased, misleading, or logically unsound reasoning patterns that appear persuasive despite being incorrect. Such issues could lead to overreliance on flawed outputs, particularly in high stakes or decision support scenarios. To mitigate these concerns, we recommend incorporating rigorous correctness verification mechanisms for both process reward models (PRMs) and reinforcement learning fine tuned models, as well as systematically auditing distilled agents for hallucinations, bias, and harmful content prior to deployment.

Looking ahead, while AgentArk has the potential to extend to real world agentic settings, such as interactive tool use workflows and safety critical decision support tasks, careful consideration of reliability, transparency, and alignment will be essential to ensure responsible and trustworthy deployment.

Table 10: Distillation results on LLaMA3-8B-Instruct. Rows denote training datasets and columns denote test benchmarks (see diagonal header). Results compare RSFT, DA, and PAD across different source models. The original LLaMA3-8B-Instruct scores 75.28 on GSM8K and 56.9 on MedMCQA.

source model	**test** \ **train**	**gsm8k**	**medmcqa**	**metamathqa**	**math**
qwen3-8b			RSFT		
	gsm8k	69.22	75.36	74.98	76.04
	medmcqa	60.08	60.39	59.86	60.87
			DA		
	gsm8k	70.89	75.06	73.24	75.82
	medmcqa	56.44	58.36	58.59	56.83
			PAD		
	gsm8k	76.42	75.28	75.55	75.72
	medmcqa	57.11	57.16	57.02	57.1
qwen3-32b			RSFT		
	gsm8k	75.77	73.54	76.15	77.18
	medmcqa	61.15	57.64	58.95	60.6
			DA		
	gsm8k	77.8	74.39	77.53	75.88
	medmcqa	56.47	58.12	58.83	56.77
			PAD		
	gsm8k	76.88	75.97	75.89	76.95
	medmcqa	60.82	60.75	60.79	60.96
gemma3-27b-it			RSFT		
	gsm8k	75.26	72.02	76.55	74.98
	medmcqa	59.77	58.76	59.5	60.94
			DA		
	gsm8k	77.89	72.33	76.37	75.59
	medmcqa	56.44	55.3	58.93	56.92
			PAD		
	gsm8k	76.35	75.97	77.02	76.27
	medmcqa	60.82	60.94	60.73	60.55

Table 11: Distillation results on Gemma-7B. Rows denote training datasets and columns denote test benchmarks (see diagonal header). Results compare RSFT, DA, and PAD across different source models. The original Gemma-7B scores 51.1 on GSM8K and 49.29 on MedMCQA.

source model	**test** \ **train**	**gsm8k**	**medmcqa**	**metamathqa**	**math**
qwen3-8b			RSFT		
	gsm8k	64.52	60.35	75.28	62.09
	medmcqa	47.17	51.37	45.04	48.51
			DA		
	gsm8k	65.73	62.02	74.75	62.62
	medmcqa	46.38	50.97	42.84	48.84
			PAD		
	gsm8k	52.74	51.69	52.34	52.55
	medmcqa	49.32	49.66	49.37	49.42
qwen3-32b			RSFT		
	gsm8k	63.23	54.66	67.25	57.62
	medmcqa	49.01	49.52	49.76	48.74
			DA		
	gsm8k	66.26	58.23	73.78	63.46
	medmcqa	48.31	48.06	48.74	49.32
			PAD		
	gsm8k	63.07	63.37	63.6	62.46
	medmcqa	49.7	50.77	49.61	50.77
gemma3-27b-it			RSFT		
	gsm8k	67.17	60.65	73.46	63.61
	medmcqa	48.31	42.03	45.66	48.46
			DA		
	gsm8k	68.16	60.27	71.34	59.59
	medmcqa	47.36	38.35	42.34	48.7
			PAD		
	gsm8k	53.37	53.15	53.24	52.62
	medmcqa	50.62	50.08	50.77	49.41

Table 12: Distillation results on Qwen3-1.7B. Rows denote training datasets and columns denote test benchmarks (see diagonal header). Results compare RSFT, DA, and PAD across different source models. The original Qwen3-1.7B scores 69.14 on GSM8K and 42.34 on MedMCQA.

source model	**test** \ **train**	**gsm8k**	**medmcqa**	**metamathqa**	**math**
qwen3-8b			RSFT		
	gsm8k	67.1	67.55	69.67	69.98
	medmcqa	43.03	42.72	43.22	42.27
			DA		
	gsm8k	68.99	69.07	69.29	69.23
	medmcqa	48.49	45.18	42.65	42.41
			PAD		
	gsm8k	69.82	69.41	69.78	69.76
	medmcqa	42.36	42.42	42.41	42.44
qwen3-32b			RSFT		
	gsm8k	69.7	65.43	69.85	68.61
	medmcqa	42.77	43.74	44.32	42.67
			DA		
	gsm8k	70.78	68.84	73.89	70.03
	medmcqa	43.68	42.46	43.03	42.72
			PAD		
	gsm8k	70.36	70.05	70.28	69.83
	medmcqa	44.41	44.46	42.43	42.46
gemma3-27b-it			RSFT		
	gsm8k	61.56	60.8	69.07	69.14
	medmcqa	42.65	39.42	41.84	42.31
			DA		
	gsm8k	71.11	65.06	69.94	64.52
	medmcqa	43.99	40.69	42.58	42.34
			PAD		
	gsm8k	69.67	70.2	69.98	69.9
	medmcqa	42.36	42.53	42.46	42.55

Table 13: Distillation results on Qwen3-0.6B. Rows denote training datasets and columns denote test benchmarks (see diagonal header). Results compare RSFT, DA, and PAD across different source models. The original Qwen3-0.6B scores 41.93 on GSM8K and 32.2 on MedMCQA.

source model	**test** \ **train**	**gsm8k**	**medmcqa**	**metamathqa**	**math**
qwen3-8b			RSFT		
	gsm8k	44.05	41.89	45.03	41.39
	medmcqa	31.08	32.9	28.19	31.7
			DA		
	gsm8k	43.67	42.23	51.78	41.17
	medmcqa	28.76	30.29	30.27	30.72
			PAD		
	gsm8k	42.57	42.25	42.51	42.53
	medmcqa	32.43	32.33	32.38	32.27
qwen3-32b			RSFT		
	gsm8k	39.04	38.51	48.14	38.74
	medmcqa	32.49	32.85	34.4	32.66
			DA		
	gsm8k	48.67	40.94	53.37	41.55
	medmcqa	34.62	32.06	34.38	34.54
			PAD		
	gsm8k	42.84	42.46	42.53	42.68
	medmcqa	33.09	34.36	34.54	32.58
gemma3-27b-it			RSFT		
	gsm8k	40.33	32.37	49.73	40.79
	medmcqa	33.28	30.5	34.4	32.01
			DA		
	gsm8k	41.02	42.46	49.2	40.94
	medmcqa	32.27	32.94	35.43	33.23
			PAD		
	gsm8k	44.61	43.52	43.72	42.17
	medmcqa	33.01	34.51	33.27	32.92

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: See details in Section 5

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

Justification: The paper does not present theoretical theorems or proof claims.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We release the code

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We release the code.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We provide the information in Section B

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Error bars are omitted due to the high computational cost of repeated experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We provide the information in Section D.3

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We offer in Section I

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: Paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: Yes, in Appendix Section H

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.