

PRIVILEGED INFORMATION DISTILLATION FOR LANGUAGE MODELS

Emiliano Penalosa^{1,2,3}, Dheeraj Vattikonda^{1,2,4}, Nicolas Gontier¹, Alexandre Lacoste¹, Laurent Charlin^{2,5}, Massimo Caccia¹

¹ServiceNow AI Research, ²Mila – Quebec AI Institute, ³Université de Montréal, ⁴McGill University, ⁵HEC Montréal

ABSTRACT

Training-time privileged information (PI) can enable language models to succeed on tasks they would otherwise fail, making it a powerful tool for reinforcement learning in hard, long-horizon settings. However, transferring capabilities learned with PI to policies that must act without it at inference time remains a fundamental challenge. We study this problem in the context of distilling frontier models for multi-turn agentic environments, where closed-source systems typically hide their internal reasoning and expose only action trajectories. This breaks standard distillation pipelines, since successful behavior is observable but the reasoning process is not. For this, we introduce π -Distill, a joint teacher-student objective that trains a PI-conditioned teacher and an unconditioned student simultaneously using the same model. Additionally, we also introduce On-Policy Self-Distillation (OPSD), an alternative approach that trains using Reinforcement Learning (RL) with a reverse KL-penalty between the student and the PI-conditioned teacher. We show that both of these algorithms effectively distill frontier agents using action-only PI. By leveraging PI as an exploration signal, our methods are substantially more data-efficient than standard RL, requiring fewer environment interactions to achieve strong performance. Specifically we find that π -Distill and in some cases OPSD, outperform industry standard practices (Supervised finetuning followed by RL) that assume access to full Chain-of-Thought supervision across multiple agentic benchmarks, models, and forms of PI. We complement our results with extensive analysis that characterizes the factors enabling effective learning with PI, focusing primarily on π -Distill and characterizing when OPSD is competitive.

1 INTRODUCTION

Language Models (LMs) have the unique ability to converse, which provides a superior user interface and more straightforward interactions than other machine learning systems. Crafting good prompts, or equivalently, conditioning on the right context, remains essential to obtain the best possible performance from an LM. This raises the question of whether LMs can learn from informative prompts to generalize to less informative prompts. In other words, how can an LM learn to transfer training-time privileged information (PI) (Vapnik & Vashist, 2009) to test tasks (that do not contain PI)?

Training-time PI can be particularly useful for Reinforcement Learning (RL) with LMs, where learning is contingent on the model’s ability to first succeed on a task (Zelikman et al., 2022; 2024). That is, the model can leverage train-time PI to succeed at tasks it would otherwise fail, effectively enabling it to bootstrap its learning from these successful experiences. As noted above, the underlying challenge is transfer: where training a model with PI and obtaining a test-time policy that retains those enhanced capabilities without PI.

In this work, we show that leveraging train-time PI is highly effective under certain conditions. For instance, conditioning a policy on PI can drastically shift the sampling distribution away from the unconditioned one, making transfer significantly harder. Additionally,

we find that non-frontier models often struggle to accurately leverage PI and must explicitly learn to use it. To enable effective training in this setting, we introduce two distillation objectives: Privileged Information Distillation (π -Distill), our main method, and On-Policy Self-Distillation (OPSD), an on-policy alternative. π -Distill adapts the typical teacher-student setup by using a single shared-parameter model in which the teacher has access to PI. Importantly, π -Distill trains the teacher and student jointly, enabling the teacher to learn how to use PI while actively mitigating distribution shift during transfer. OPSD similarly allows the teacher and student to share parameters, regularizing on-policy RL with a reverse KL penalty between the student and a PI-conditioned teacher.

We ground our work in the task of distilling frontier models for complex multi-turn agentic settings. Typically, the industry standard for these tasks involves Supervised Fine-Tuning (SFT) on frontier model outputs followed by Reinforcement Learning (RL). Unfortunately, it has become common practice for closed-source frontier models to hide their full Chain-of-Thought (CoT) reasoning, at best producing a summary alongside the action they intend to take (OpenAI et al., 2024). This opacity undermines standard distillation methods, as we can observe what successful agents do but not their reasoning.

We find that π -Distill is highly effective at mitigating the lack of CoT, outperforming industry standards that assume full CoT access, with this holding for OPSD in some cases. Moreover, by using PI to guide exploration, both methods are significantly more data-efficient than standard RL, requiring fewer environment interactions to achieve strong performance. We demonstrate this on two agentic tool-use environments, Travel Planner (Xie et al., 2024) and τ -Bench retail (Yao et al., 2024), while showing proficient Out-of-Domain generalization on the 7 tool-use environments provided by GEM (Liu et al., 2025b) and τ -Bench airline. Our findings are consistent across three models from two distinct families.

Finally, we transform frontier model trajectories into three varying types of PI, each with different information density, providing varying amounts of utility and inducing different degrees of distributional shift between student and teacher. We use this variation to analyze the critical factors for training with PI, finding that for π -Distill, maximizing PI utility while simultaneously mitigating the student-teacher distributional gap and preventing collapse is essential for effective learning. While information content is more important for OPSD.

2 BACKGROUND

3 AGENTIC INTERACTION AND POLICY OPTIMIZATION

We formalize long-horizon, multi-turn agentic environments as a Markov Decision Process (MDP). In this setting, a policy $\pi_\theta(\cdot | \mathbf{s})$ interacts with an environment over extended sequences of actions. To simplify notation, we let \mathbf{s} represent the *evolving interaction context*, which aggregates all information available to the model at a given point: the initial user prompt, the models past outputs, and all environment responses. As the agent acts, the context updates via a transition function $\mathbf{s}_{t+1} \sim P(\cdot | \mathbf{s}_t, \mathbf{o}_t)$, where the environment appends its response to the current sequence to form the next state. As \mathbf{s} encapsulates the interaction history, we use the following to simplify notation: $\pi_\theta(\mathbf{o} | \mathbf{s}) = \prod_{i=0}^T \pi_\theta(\mathbf{z}_i, \mathbf{a}_i | \mathbf{s}_{<i})$ where $\mathbf{o} = (\mathbf{z}, \mathbf{a})$ consists of reasoning tokens \mathbf{z} and action tokens \mathbf{a} . RL seeks to maximize the expected return obtained through interaction with an environment (Sutton et al., 1999): $J(\pi_\theta) = \mathbb{E}_{\mathbf{o} \sim \pi_\theta(\cdot | \mathbf{s})} [R(\mathbf{o}, \mathbf{s})]$. Here, $R(\mathbf{o}, \mathbf{s}) \in [-1, 1]$ denotes the reward assigned by the environment to the generated trajectory. We optimize

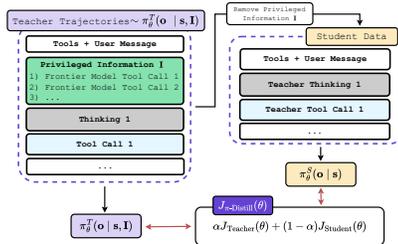


Figure 1: **Overview of the π -Distill framework.** (1) Successful trajectories (not shown) are collected from a frontier agent that exposes only actions while hiding its Chain-of-Thought. (2) These trajectories are transformed into training-time privileged information (PI) and used to sample a PI-conditioned teacher policy $\pi_\theta^T(\mathbf{o} | \mathbf{s}, \mathbf{I})$. (3) The PI-conditioned teacher and an unconditioned student $\pi_\theta^S(\mathbf{o} | \mathbf{s})$ share parameters and are trained jointly, enabling transfer of privileged knowledge to a test-time policy that acts without PI.

$J(\pi_\theta)$ using Group Relative Policy Optimization (GRPO) (Shao et al., 2024; DeepSeek-AI et al., 2025), with the adjustments recommended by Yu et al. (2025) and Liu et al. (2025a). For each state \mathbf{s}_i , we sample a group of G trajectories $\{\mathbf{o}_g\}_{g=1}^G$ according to the current sampling policy μ and the transition function P . For each token k in trajectory g , we define the token-level importance ratio: $\rho_{g,k}(\theta) = \frac{\pi_\theta(o_{g,k}|\mathbf{s}_i, \mathbf{o}_{g,<k})}{\mu(o_{g,k}|\mathbf{s}_i, \mathbf{o}_{g,<k})}$. We define a group-relative advantage $A_{s,g}$ by comparing the return of trajectory g to the average return of the sampled group, and use it to scale clipped importance-weighted policy updates. The GRPO objective then is:

$$J_{\text{GRPO}}(\theta) = \mathbb{E}_{\substack{\mathbf{o} \sim \pi_\theta(\cdot|\mathbf{s}) \\ \mathbf{s} \sim P}} \left[\frac{1}{\sum_g K_g} \sum_{g,k} \min \left(\rho_{g,k} A_{s,g}, \text{clip}(\rho_{g,k}, 1-\epsilon, 1+\epsilon) A_{s,g} \right) \right] \quad (1)$$

where K_g is the length of trajectory g in tokens and ϵ is the clipping parameter. While the original GRPO objective typically employs a KL-penalty with respect to the base model, we drop this term as recent work shows it can hinder performance (Shah et al., 2026).

4 METHODS

In this section, we introduce and motivate two algorithms for leveraging PI during training. Both methods use a teacher-student framework inspired by traditional distillation. However, the teacher and student share parameters, and only the teacher is conditioned on PI.

π -Distill learns from teacher-generated traces by jointly improving both the teacher and the student. In contrast, on-policy self-distillation (OPSD)¹ samples trajectories from the student policy and uses a reverse KL divergence between the student and the PI-conditioned teacher as a training penalty.

4.1 PRIVILEGED INFORMATION DISTILLATION (π -DISTILL)

Motivation A straightforward way to use PI for distillation is to condition a policy on PI to generate successful trajectories, which are then used for fine-tuning. However, this approach has two key limitations. First, base models do not automatically know how to exploit PI (see section 7), needing to learn how to use it, before providing benefit. Second, even once the policy can exploit PI, we still need to transfer this behavior to a policy that must act without PI at test time. A naive solution is to first train a PI-conditioned policy and then distill its behavior into an unconditioned one. In practice, this sequential pipeline introduces several issues. It is unclear which checkpoint of the conditioned policy should be distilled, learning from its trajectories is off-policy and can be unstable, and training the two policies separately is computationally inefficient. Our early experiments confirm that this setup leads to suboptimal performance (see fig. 4). To address these challenges, we propose Privileged Information Distillation (π -DISTILL), which trains both policies jointly within a single parameter-shared model. This allows the model to learn to exploit PI while simultaneously learning to act without it.

Algorithm. Our approach uses a single model with shared parameters θ that acts as both a teacher $\pi_\theta^T(\mathbf{o} | \mathbf{s}, \mathbf{I})$ (conditioned on PI \mathbf{I}) and a student $\pi_\theta^S(\mathbf{o} | \mathbf{s})$ (operating without PI). We train both simultaneously using two objectives. The teacher objective trains the conditioned policy to maximize reward while maintaining proximity to the student policy:

$$J_{\text{Teacher}}(\theta) = \mathbb{E}_{\substack{\mathbf{o} \sim \pi_\theta^T(\mathbf{o}|\mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} \left[R(\mathbf{o}, \mathbf{s}) \right] - \beta D_{\text{KL}} \left(\pi_\theta^T(\mathbf{o} | \mathbf{s}, \mathbf{I}) \parallel \text{sg} \left(\pi_\theta^S(\mathbf{o} | \mathbf{s}) \right) \right). \quad (2)$$

This objective samples trajectories from the teacher policy and updates it to increase reward while subject to a reverse KL penalty $D_{\text{KL}}(\pi_\theta^T \parallel \pi_\theta^S)$ controlled by β , where $\text{sg}(\cdot)$ indicates the stop gradient operator. This objective serves two purposes: (i) it encourages the teacher to fit high-reward modes familiar to the student, making learning from its traces easier,

¹We adopt the naming "self-distillation" following concurrent work Shenfeld et al. (2026); Zhao et al. (2026) that propose the same objective.

and (ii) shared parameters promote transfer of the teachers knowledge to the student, even without directly training the student.

The student objective trains the unconditioned policy to learn from the teacher’s trajectories:

$$J_{\text{Student}}(\theta) = \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | s, \mathbf{I}) \\ \mathbf{s} \sim P}} \left[\frac{\pi_{\theta}^{\text{S}}(\mathbf{o} | s)}{\text{sg}(\pi_{\theta}^{\text{T}}(\mathbf{o} | s, \mathbf{I}))} R(\mathbf{o}, \mathbf{s}) \right] - \beta D_{\text{KL}}(\text{sg}(\pi_{\theta}^{\text{T}}(\mathbf{o} | s, \mathbf{I})) \parallel \pi_{\theta}^{\text{S}}(\mathbf{o} | s)) \quad (3)$$

This objective samples trajectories from the teacher (which has access to PI) but updates the student policy (which does not). This teaches the student to replicate the teacher’s high-reward behavior without needing PI.

Combining the student and teacher terms gives us our final objective:

$$J_{\pi\text{-Distill}}(\theta) = \alpha J_{\text{Teacher}}(\theta) + (1 - \alpha) J_{\text{Student}}(\theta), \quad (4)$$

where $\alpha \in [0, 1]$ controls the balance between student and teacher focused learning.

When $\alpha = 1$, optimization focuses entirely on the teacher, although the student may still improve through shared parameters. When $\alpha = 0$, training is focused on student learning from the teachers current behavior. Where, we observe that under certain conditions, parameter sharing can still lead to improvements in the teacher without explicit teacher updates. When $\alpha = 0.5$, both are optimized jointly. Shared parameters allow representations learned for using PI to transfer to the student, while student updates keep those representations effective without PI. The full algorithm is given in algorithm 1.

Connection to Variational EM. This approach can be viewed as a form of Variational Expectation-Maximization (EM), where one uses an approximate posterior π^{T} to approximate a target distribution π^* . Here the E-step first improves π^{T} and the M-step distills this into the student policy π^{S} . Traditionally this can be trained sequentially or in alternating loops with separate models (see Zhou et al. (2025)). We discuss this connection more in depth in App. C.1 and App. C.2, characterizing the target distribution π^* , derive $\pi\text{-Distill}$ from this perspective and compare against sequential setups (similar to Zhou et al. (2025)).

4.2 ON-POLICY SELF DISTILLATION

Motivation. $\pi\text{-Distill}$ can be viewed as off-policy learning, where the student is trained on trajectories generated by the PI-conditioned teacher. A complementary line of work studies on-policy distillation, in which the student acts as the sampling policy and knowledge is transferred by minimizing the reverse KL between the student and teacher (Agarwal et al., 2024). Prior work typically utilizes a larger model as the teacher. We introduce this objective in our PI setting by instantiating it with the same shared-parameter model, where the teacher is additionally conditioned on PI. We refer to this objective as On-Policy Self-Distillation (OPSD) to maintain consistency with concurrent work (Shenfeld et al., 2026; Zhao et al., 2026; Hübötter et al., 2026).

Algorithm. The above intuition yields the following objective:

$$J_{\text{OPSD}}(\theta) = \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} | s) \\ \mathbf{s} \sim P}} \left[R(\mathbf{o}, \mathbf{s}) \right] - \beta D_{\text{KL}}(\pi_{\theta}^{\text{S}}(\mathbf{o} | s) \parallel \text{sg}(\pi_{\theta}^{\text{T}}(\mathbf{o} | s, \mathbf{I}))). \quad (5)$$

Note that the updates are on-policy as the expectation is taken over π^{S} . Where the reverse KL acts as a dense per-token reward measuring how closely the student matches the teacher. The full algorithm is given in algorithm 2. We analyze OPSD in greater depth in App D and characterize the specific target distribution that the algorithm implicitly fits. Concurrent work also propose this objective, demonstrating its effectiveness in settings such as having access to ground truth answers (Zhao et al., 2026; Hübötter et al., 2026) and continual learning (Shenfeld et al., 2026). In this work, we introduce and evaluate OPSD for PI transfer without ground truth, and identify the settings in which it fails.

5 EXPERIMENTAL SETTING

While learning with PI is applicable to many settings, we ground our work in the task of distilling frontier models within multi-turn tool-calling environments. We focus on this domain for two primary reasons. First, non-frontier models often lack the capabilities required for such settings, whereas frontier models demonstrate proficient performance (Singh et al., 2025; Team et al., 2025). Consequently, weaker models struggle to independently sample successful trajectories, making the PI derived from a frontier model’s actions highly valuable. Second, frontier models typically occlude their CoT reasoning (OpenAI et al., 2024). This renders standard distillation methods infeasible, a gap our proposed algorithms are designed to fill. We outline benchmarks, sources of PI, models and benchmarks here, deferring implementation details to App appendix J.

5.1 BENCHMARKS

To evaluate our approach, we first employ τ -Bench (Yao et al., 2024), which simulates customer service interactions where agents book flights (airline domain) or assist shoppers (retail domain) by calling tools and gathering user information. To reduce computational costs, we substitute the GPT-4o user simulator with Qwen-14B. The most substantial change is that we remove the `transfer_to_human_agents` tool, as it consistently led to reward hacking. The resulting dataset consists of 500 training tasks restricted to the retail domain; we evaluate on 115 held-out retail tasks and 50 airline tasks, utilizing the latter to test Out-Of-Domain (OOD) generalization. Next, we utilize Travel Planner (Xie et al., 2024), a benchmark focusing on tool use for planning. While the original repository² employs a rubric-based evaluation that prioritizes "easy" constraints before checking "hard" ones, we found this setup causes policies to collapse onto undesired behaviors. To address this, we decouple the rewards so that easy constraints are tied directly to their corresponding hard constraints, for example, verifying dietary restrictions immediately after booking a restaurant. We train on the 45 training tasks and report results on the 180 publicly available held-out tasks. Finally, to probe whether training on these environments enhances tool usage to scopes far beyond the training domains, we evaluate on the GEM QA multi-turn tool-usage environment suite (Liu et al., 2025b). This suite equips agents with a search tool consisting of seven environments (2Wiki (Ho et al., 2020), PopQA (Mallen et al., 2022), TriviaQA (Joshi et al., 2017), HotpotQA (Yang et al., 2018), Bamboogle (Press et al., 2023), NaturalQuestions (Kwiatkowski et al., 2019), and Musique (Trivedi et al., 2022)). See App. G for experiments on OOD settings.

5.2 SOURCES OF PRIVILEGED INFORMATION

To show the benefits of leveraging train-time PI, we aim to distill a frontier model using only its raw output trajectories. For this, we mine trajectories from `DeepSeek-chat-v3.1` as it is open-source and allows access to its reasoning tokens. Having access to these tokens enables us to benchmark against standard settings that assume access to CoT, enabling us to analyze the performance impact when one lacks access to the expert’s reasoning traces.

Types of Privileged Information. We transform the raw trajectories into three distinct forms of PI to analyze how varying information density impacts performance: **Tool calls & arguments**, which include the complete sequence of actions executed by the frontier model, including function names and input arguments (e.g., `GetUserDetails(Name:"Kevin Lau")`); **Tool calls only**, where we provide only the function names such as `GetUserDetails` and require the model to infer the correct arguments from context; and **Self-generated hints**, where we prompt the *trained* model to summarize a successful trajectory produced by the frontier model. This process can help filter inefficient actions and suggest ways to improve upon the experts behavior (see App. L.3 for an example).

Leveraging these three formats enables us to identify what properties allow successful transfer between teacher and student. We analyze this in section 7 finding that the optimal

²<https://github.com/OSU-NLP-Group/TravelPlanner>

Table 1: **Evaluation results on Travel Planner, τ -Bench (Retail), and τ -Bench (Airline)**. Shaded rows denote our methods. **Bold** values indicate the best performance within each model category, while underlined values indicate the second-best. Results show mean \pm standard deviation across three random seeds.

	Travel Planner	τ -Bench Retail	τ -Bench Airline (OOD)
<i>DeepSeek V3.1 Chat-671B</i>			
Base	45.0% \pm 3.78	51.3% \pm 0.212	40.0% \pm 0.161
<i>R1-Distill-Llama-8B</i>			
Base	0.00% \pm 0.00	0.00% \pm 0.00	0.00% \pm 0.00
SFT w/ CoT	6.35% \pm 1.27	15.2% \pm 0.44	8.00% \pm 5.29
SFT w/o CoT	2.40% \pm 0.86	0.680% \pm 0.17	0.670% \pm 1.15
SFT w/ CoT + RL	12.4% \pm 1.56	16.3% \pm 1.49	7.33% \pm 4.16
SFT w/ CoT + On-Policy Self Distillation	13.1% \pm 1.09	14.5% \pm 0.00	2.00% \pm 0.00
SFT w/ CoT + π -Distill π^S ($\alpha = 0$)	7.86% \pm 1.75	18.6% \pm 0.50	10.0% \pm 2.00
SFT w/ CoT + π -Distill $\pi^S + \pi^T$ ($\alpha = 0.5$)	<u>14.0%</u> \pm 1.63	<u>18.3%</u> \pm 0.77	<u>9.33%</u> \pm 3.06
SFT w/ CoT + π -Distill π^T ($\alpha = 1$)	14.1% \pm 3.27	17.7% \pm 0.77	7.33% \pm 4.16
<i>QWEN3-4B</i>			
Base	17.6% \pm 2.16	5.03% \pm 1.88	2.21% \pm 1.99
SFT w/ CoT	21.1% \pm 1.94	12.4% \pm 0.60	2.67% \pm 1.15
SFT w/o CoT	20.8% \pm 1.82	15.2% \pm 1.70	4.00% \pm 5.29
RL	25.1% \pm 3.75	15.2% \pm 0.17	5.33% \pm 3.06
SFT w/o CoT + RL	23.3% \pm 1.33	17.6% \pm 2.69	5.33% \pm 2.31
On-Policy Self Distillation	29.8% \pm 1.14	23.1% \pm 0.04	<u>10.6%</u> \pm 6.57
π -Distill π^S ($\alpha = 0$)	28.5% \pm 4.35	25.3% \pm 0.60	8.00% \pm 5.29
π -Distill $\pi^S + \pi^T$ ($\alpha = 0.5$)	33.8% \pm 6.85	22.6% \pm 0.93	6.00% \pm 2.00
π -Distill π^T ($\alpha = 1$)	28.2% \pm 6.27	22.5% \pm 0.93	12.0% \pm 5.29
SFT w/ CoT + RL	26.4% \pm 1.16	<u>23.3%</u> \pm 3.02	6.67% \pm 5.77
<i>QWEN3-8B</i>			
Base	23.6% \pm 2.23	3.35% \pm 1.47	6.40% \pm 3.02
SFT w/ CoT	26.0% \pm 2.27	16.5% \pm 4.66	5.33% \pm 1.15
SFT w/o CoT	29.8% \pm 1.71	12.8% \pm 0.77	6.00% \pm 4.00
RL	27.3% \pm 0.95	23.9% \pm 0.44	6.67% \pm 3.06
SFT w/o CoT + RL	31.3% \pm 2.79	23.5% \pm 4.27	6.00% \pm 2.00
On-Policy Self Distillation	37.5% \pm 1.53	27.3% \pm 0.33	14.0% \pm 5.66
π -Distill π^S ($\alpha = 0$)	40.7% \pm 1.14	31.1% \pm 0.73	<u>12.0%</u> \pm 6.00
π -Distill $\pi^S + \pi^T$ ($\alpha = 0.5$)	<u>41.1%</u> \pm 7.24	<u>30.6%</u> \pm 0.67	7.33% \pm 1.15
π -Distill π^T ($\alpha = 1$)	44.1% \pm 2.16	29.7% \pm 0.33	9.33% \pm 3.06
SFT w/ CoT + RL	32.3% \pm 1.10	29.1% \pm 2.14	8.00% \pm 3.46

configuration depends on the value of α and the specific attributes of the PI (e.g. KL between teacher and student).

Using this setup, we collect 15,885 successful traces for τ -Bench retail and 1,986 for Travel Planner, evenly sampled across tasks. From these traces, for each training task we select the successful trajectory with the least number of steps to build the PI. In total we obtain PI for all 45 tasks in Travel Planner and 300/500 tasks in τ -Bench. We note we exclude the self-generated hints for R1-Deepseek-Distill-LLama-8B as it consistently returns the raw input trace or tool calls as the hint.

5.3 MODELS AND BASELINES

Models. We employ Qwen3-4B and Qwen3-8B (Yang et al., 2025) being strong reasoning models. We also evaluate R1-Distill-Llama-8B to cover a distinct model family. We find this model fails to generate correct trajectories even when conditioned on PI in both benchmarks, making direct RL training unfeasible. Thus, exclusively for R1-Distill-Llama-8B, we warm-start it using SFT w/ CoT from expert traces. This setup allows us to determine if PI remains beneficial even when the model has already seen it during training.

Baselines. We instantiate π -Distill with $\alpha \in \{0, 0.5, 1\}$ and OPSD. We compare against (i) standard RL, (ii) SFT on expert trajectories with and without CoT, and (iii) SFT followed by RL. We utilize the GRPO objective outlined in eq. (1) for all RL-based algorithms. Following Vattikonda et al. (2025), we sweep over multiple SFT checkpoints for SFT+RL baselines and report results using the checkpoint that yields the strongest final performance. We use the full set of collected successful traces for all SFT baselines, as this maximized performance. For π -Distill and OPSD in τ -Bench, we utilize PI whenever available, otherwise, we sample the traces with the student and perform regular RL for that goal.

6 MAIN RESULTS

In this section, we demonstrate that leveraging PI provides a potent learning signal on held-out tasks, proving effective even when full CoT supervision is available. Table 1 details performance metrics for Travel Planner and τ -Bench. Our primary finding is that π -Distill

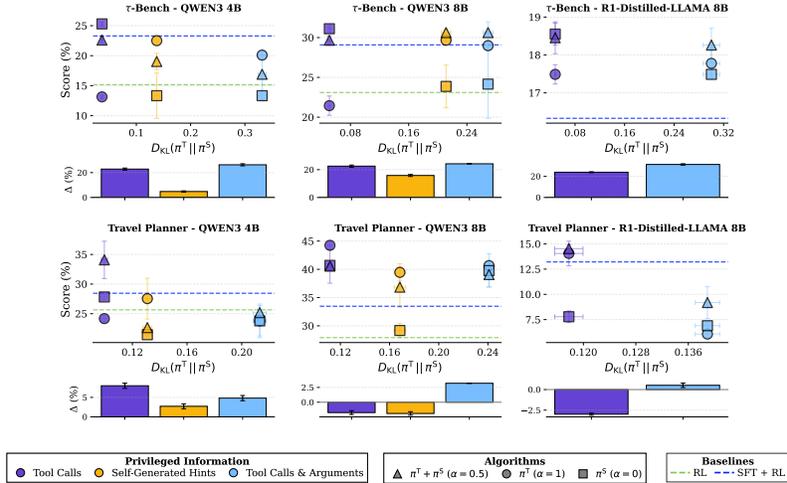


Figure 2: **Impact of PI Types and Algorithms on Performance.** We compare held-out performance on τ -Bench (top row) and Travel Planner (bottom row) across three base models and three PI types (colors). The scatter plots map final scores against the initial teacher-student divergence ($D_{\text{KL}}(\pi_{\text{base}}^T \parallel \pi_{\text{base}}^S)$), while the bar-charts display the PI utility (Δ) on training tasks.

variants achieve superior performance across all but a single setting. With Qwen3-4B we find that π -Distill using $\alpha = 0$ can outperform SFT + RL w/ CoT, with other variants approximating its performance. Notably, in Qwen3-8B, π -Distill consistently outperforms the industry standard (SFT w/ CoT + RL) regardless of α . In the best-case scenarios, π -Distill achieves substantial improvements: 11.8% on Travel Planner, and 2.08% and 6.00% on the retail and airline subsets of τ -Bench, respectively. These results are substantial as they confirm that π -Distill effectively distills frontier models even when CoT traces are hidden, enabling non-frontier models to become proficient at complex multi-step agentic tasks. Moreover, π -Distill achieves this with significantly greater efficiency. Unlike SFT w/ CoT + RL, which requires sweeping over multiple SFT checkpoints to achieve peak performance (Vattikonda et al., 2025), π -Distill requires only a single training phase, greatly simplifying the training process.

Regarding α , we observe that there is no definitive best value, rather the best-performing α varies by setting. We analyze this nuance further in § 7. When analyzing other methods that do not assume access to full CoT, we find RL and SFT w/o CoT + RL perform similarly, while RL requires significantly less compute, but find both fail to yield substantial gains, significantly lagging behind SFT w/ CoT + RL. Additionally, we find that OPSD, can substantially outperform these baselines in τ -Bench, while performing similarly in Travel Planner. We also observe that when OPSD succeeds, it provides substantial gains in OOD settings (τ -Bench Airline), scaling with model capacity, being the second best in Qwen3-4B and best in Qwen3-8B. Finally, our results on R1-Distill-Llama-8B show that π -Distill remains useful even if the model has previously been SFTd on traces containing the PI. We defer further OOD experiments to App G.

7 WHAT MATTERS WHEN USING TRAIN-TIME PI

While prior sections focused on the best-performing configurations, here we analyze how varying the type of PI affects final performance. Our goal is to isolate the factors that determine success when training with PI. We identify two primary drivers: (i) the divergence between the conditioned and unconditioned base policies ($D_{\text{KL}}(\pi_{\text{base}}^T \parallel \pi_{\text{base}}^S)$ for π -Distill and $D_{\text{KL}}(\pi_{\text{base}}^S \parallel \pi_{\text{base}}^T)$ for OPSD), and (ii) the usefulness of the privileged signal, captured by the utility $\Delta = \text{score}(\pi_{\text{base}}^T) - \text{score}(\pi_{\text{base}}^S)$ on *training tasks*. Additionally, in fig. 7 we report the maximum attainable utility on $\Delta_{\text{max}} = \max_t \text{score}(\pi_t^{\text{PI}}) - \max_t \text{score}(\pi_t^{\text{RL}})$, defined as the difference between the best scores on *training tasks* achieved with PI and

without PI (pure RL), which measures how effectively each algorithm converts access to PI into performance gains. We defer OPSD analysis to App. G.1

7.1 WHAT MATTERS FOR π -DISTILL

Figure 2 displays held-out performance on both τ -Bench Retail and Travel Planner across the types of PI defined in § 5.2. Each subplot maps final performance against the initial divergence $D_{\text{KL}}(\pi_{\text{base}}^{\text{T}} \parallel \pi_{\text{base}}^{\text{S}})$, with accompanying bar plots showing the utility Δ for each type of PI. The most prominent pattern is that as the initial KL increases, final performance tends to decrease, though the full picture requires a more nuanced, α -dependent analysis.

Teacher only training $\alpha = 1$. In this setting, we find performance generally declines or maintains as KL divergence increases. The primary exception occurs when using only tool calls in τ -Bench, where this variant underperforms. We trace this failure mode to an early collapse in KL divergence. Figure 6 displays the KL as training progresses. We observe that that even when $\beta = 0$ the KL drops to near zero, indicating teacher and student have collapsed onto each other ($\pi^{\text{T}} \approx \pi^{\text{S}}$). As a consequence we find the teacher π^{T} learns to ignore the PI causing it to underperform even the RL baseline. We attempted to mitigate this by using π_{base} instead of π^{S} when calculating the KL, but found this further degraded results (see App. H.2). In addition, our results confirm that training the teacher is an important part of leveraging PI. We can see this from the observation that even when the initial utility is negative ($\Delta < 0$), training the teacher allows the policy to learn to leverage the PI, consistently showing positive Δ_{max} values in such cases. We find that effectively learning to use the PI is a significant contributing factor to transfer when only training the teacher.

Joint training $\alpha = 0.5$. We identify $\alpha = 0.5$ as the most robust configuration. It achieves the best performance in 7 out of 16 scenarios and effectively avoids the failure modes of the other variants, ranking as the worst performer only once. By balancing both teacher and student objectives, $\alpha = 0.5$ is able to mitigate the failure cases of independent training. We believe more granular tuning of α can likely lead to optimal performance in most settings, leaving this as future work.

Student only training $\alpha = 0$. Here, low KL divergence is a strong predictor of success. In τ -Bench, for example, setting $\alpha = 0$ with only tool calls consistently yields the best results, as the minimal distribution shift makes learning from the conditioned traces significantly easier. As KL divergence rises, performance generally drops, though high-utility PI can occasionally reverse this trend. Further, we find that PI utility Δ can play a large role in the success of this variant. For instance, on Travel Planner with **Qwen3-8B**, $\alpha = 0$ underperforms significantly because the PI provides negative utility ($\Delta < 0$), actively degrading performance. Furthermore, as seen in fig. 7, we find that $\Delta_{\text{max}} \approx 0$ in these failure cases, confirming that the algorithm cannot extract value from the PI when the signal itself offers no initial advantage over the base model. On the other hand, we find that when the teacher-student KL is low, student-only training effectively transfers knowledge back onto the teacher. This allows the teacher to learn how to leverage PI even though it is not being directly trained. We observe this behavior in **Qwen3-8B** on Travel Planner when using only tool calls as PI. While the initial $\Delta < 0$, the low KL allows transfer to occur from the student to the teacher and enables the teacher to sample improved traces.

8 CONCLUSION

We introduce π -Distill and On-Policy Self-Distillation, two algorithms that leverage PI at training time to produce an improved policy even when that information is lacking at test time. By grounding our work in the distillation of frontier models where CoT reasoning is inaccessible, we demonstrate that the absence of these proprietary traces is not a limiting factor. In fact, our approach remains effective even when compared to industry standards that assume access to full CoT reasoning. Through evaluations across two training domains and eight OOD datasets, we find that π -Distill and OPSD are highly effective in all explored

settings. Finally, we characterize the factors driving the distillation of training-time PI, showing that its success can often be predicted using only base-model statistics.

9 ACKNOWLEDGMENTS

We thank and credit Siddarth Venkatraman for the idea to view learning from PI from a variational perspective, which led to the final version of π -Distill. We thank Vedant Shah for helpful discussion and proof reading initial drafts of the paper. Additionally, we thank Michael Noukhovitch for valuable feedback that substantially aided the final draft of this work. We also thank Alexandre Piché for their valuable discussions.

EP acknowledges the support of the NSERC PGS-D grant, the Bourse en intelligence artificielle provided by Université de Montréal and the MITACS Accelerate grant. DV acknowledges the support of the MITACS Accelerate grant. LC recognizes the support of NSERC, the Canada CIFAR AI Chair Program, the Canada First Research Excellence Fund and IVADO.

10 AUTHOR CONTRIBUTIONS

Emiliano Penalzoa proposed the idea of leveraging PI, led the project, implemented all algorithms and the RL environments, conducted most experiments, and was the primary writer of the paper.

Dheeraj Vattikonda was a core part of the project, with the final version heavily improved by their contributions. Specifically, they provided substantial support on all experiments over the course of the project, implemented and conducted experiments on OOD environments, and provided heavy contributions to paper writing and figures.

Nicolas Gontier helped in advising the project and providing detailed feedback that substantially improved experimental rigor.

Alexandre Lacoste advised the project, providing valuable insights to the goal and purpose of leveraging PI, helped improve experimental rigor and analysis, and assisted in paper writing.

Laurent Charlin advised the project, helping influence the framing and perspective of the project, and provided help with writing, with a substantial portion of the introduction written by them.

Massimo Caccia was a core member of the project. They helped shape the research direction, proposed the on-policy self-distillation component, guided the design of extensive ablation studies, and provided detailed technical guidance across all stages. They also contributed substantially to writing and figures.

REFERENCES

- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes, 2024. URL <https://arxiv.org/abs/2306.13649>.
- Lucas Caccia, Alan Ansell, Edoardo Ponti, Ivan Vuli, and Alessandro Sordoni. Training plug-n-play knowledge modules with deep context distillation, 2025. URL <https://arxiv.org/abs/2503.08727>.
- Bowen Cao, Deng Cai, and Wai Lam. Infiniteicl: Breaking the limit of context window size via long short-term memory transformation, 2025. URL <https://arxiv.org/abs/2504.01707>.
- Justin Chih-Yao Chen, Becky Xiangyu Peng, Prafulla Kumar Choubey, Kung-Hsiang Huang, Jiaxin Zhang, Mohit Bansal, and Chien-Sheng Wu. Nudging the boundaries of llm reasoning, 2025. URL <https://arxiv.org/abs/2509.25666>.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhen Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. URL <https://arxiv.org/abs/2501.12948>.

Aniket Didolkar, Anirudh Goyal, Nan Rosemary Ke, Siyuan Guo, Michal Valko, Timothy Lillicrap, Danilo Rezende, Yoshua Bengio, Michael Mozer, and Sanjeev Arora. Metacognitive capabilities of llms: An exploration in mathematical problem solving, 2024. URL <https://arxiv.org/abs/2405.12205>.

Aniket Didolkar, Nicolas Ballas, Sanjeev Arora, and Anirudh Goyal. Metacognitive reuse: Turning recurring llm reasoning into concise behaviors, 2025. URL <https://arxiv.org/abs/2509.13237>.

Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6609–6625, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.coling-main.580>.

Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, and Nikolay Malkin. Amortizing intractable inference in large language models, 2024. URL <https://arxiv.org/abs/2310.04363>.

Yukun Huang, Yanda Chen, Zhou Yu, and Kathleen McKeown. In-context learning distillation: Transferring few-shot learning ability of pre-trained language models, 2022. URL <https://arxiv.org/abs/2212.10670>.

Jonas Hübner, Frederike Lübeck, Lejs Behric, Anton Baumann, Marco Bagatella, Daniel Marta, Ido Hakimi, Idan Shenfeld, Thomas Kleine Büning, Carlos Guestrin, and Andreas Krause. Reinforcement learning via self-distillation, 2026. URL <https://arxiv.org/abs/2601.20802>.

- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 2017.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Chengpeng Li, Mingfeng Xue, Zhenru Zhang, Jiayi Yang, Beichen Zhang, Xiang Wang, Bowen Yu, Binyuan Hui, Junyang Lin, and Dayiheng Liu. Start: Self-taught reasoner with tools, 2025. URL <https://arxiv.org/abs/2503.04625>.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective, 2025a. URL <https://arxiv.org/abs/2503.20783>.
- Zichen Liu, Anya Sims, Keyu Duan, Changyu Chen, Simon Yu, Xiangxin Zhou, Haotian Xu, Shaopan Xiong, Bo Liu, Chenmian Tan, Chuen Yang Beh, Weixun Wang, Hao Zhu, Weiyan Shi, Diyi Yang, Michael Shieh, Yee Whye Teh, Wee Sun Lee, and Min Lin. Gem: A gym for agentic llms, 2025b. URL <https://arxiv.org/abs/2510.01051>.
- Renjie Luo, Zichen Liu, Xiangyan Liu, Chao Du, Min Lin, Wenhui Chen, Wei Lu, and Tianyu Pang. Language models can learn from verbal feedback without scalar rewards, 2025. URL <https://arxiv.org/abs/2509.22638>.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Hannaneh Hajishirzi, and Daniel Khashabi. When not to trust language models: Investigating effectiveness and limitations of parametric and non-parametric memories. *arXiv preprint*, 2022.
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich, Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghorbani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi, Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts, Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong, Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang, Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpouras, Francis Song, Fred von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Haiming Bao, Hao Sheng, Hart Andrin, Hessam Bagherinezhad, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ian Kivlichan, Ian O’Connell, Ian Osband, Ignasi Clavera Gilaberte, Ilge Akkaya, Ilya Kostrikov, Ilya Sutskever, Irina Kofman, Jakub Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñero Candela, Joe Palermo, Joel Parish, Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus, Lillian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondraciuk, Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko Tintor, Mason Meyer, Matt Jones, Matt Kaufer, Max Schwarzer, Meghan Shah, Mehmet Yatbaz, Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe, Michael Malek, Michele Wang, Michelle Fradin,

- Mike McClay, Mikhail Pavlov, Miles Wang, Mingxuan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowdhury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias, Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng, Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie Qi, Vineet Kosaraju, Vinnie Monaco, Vitthyr Pong, Vlad Fomenko, Weiyi Zheng, Wenda Zhou, Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yining Chen, Young Cha, Yu Bai, Yuchen He, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. Openai o1 system card, 2024. URL <https://arxiv.org/abs/2412.16720>.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. Measuring and narrowing the compositionality gap in language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 5687–5711, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.378. URL <https://aclanthology.org/2023.findings-emnlp.378/>.
- Yuxiao Qu, Amrith Setlur, Virginia Smith, Ruslan Salakhutdinov, and Aviral Kumar. Pope: Learning to reason on hard problems via privileged on-policy exploration, 2026. URL <https://arxiv.org/abs/2601.18779>.
- Vedant Shah, Johan Obando-Ceron, Vineet Jain, Brian Bartoldson, Bhavya Kailkhura, Sarthak Mittal, Glen Berseth, Pablo Samuel Castro, Yoshua Bengio, Nikolay Malkin, Moksh Jain, Siddarth Venkatraman, and Aaron Courville. A comedy of estimators: On kl regularization in rl training of llms, 2026. URL <https://arxiv.org/abs/2512.21852>.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.
- Idan Shenfeld, Mehul Damani, Jonas Hübötter, and Pulkit Agrawal. Self-distillation enables continual learning, 2026. URL <https://arxiv.org/abs/2601.19897>.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, Akshay Nathan, Alan Luo, Alec Helyar, Aleksander Madry, Aleksandr Efremov, Aleksandra Spyra, Alex Baker-Whitcomb, Alex Beutel, Alex Karpenko, Alex Makelov, Alex Neitz, Alex Wei, Alexandra Barr, Alexandre Kirchmeyer, Alexey Ivanov, Alexi Christakis, Alistair Gillespie, Allison Tam, Ally Bennett, Alvin Wan, Alyssa Huang, Amy McDonald Sandjideh, Amy Yang, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrei Gheorghe, Andres Garcia Garcia, Andrew Braunstein, Andrew Liu, Andrew Schmidt, Andrey Mereskin, Andrey Mishchenko, Andy Applebaum, Andy Rogerson, Ann Rajan, Annie Wei, Anoop Kotha, Anubha Srivastava, Anushree Agrawal, Arun Vijayvergiya, Ashley Tyra, Ashvin Nair, Avi Nayak, Ben Eggers, Bessie Ji, Beth Hoover, Bill Chen, Blair Chen, Boaz Barak, Borys Minaiev, Botao Hao, Bowen Baker, Brad Lightcap, Brandon McKinzie, Brandon Wang, Brendan Quinn, Brian Fioca, Brian Hsu, Brian Yang, Brian Yu, Brian Zhang, Brittany Brenner, Callie Riggins Zetino, Cameron Raymond, Camillo Lugaresi, Carolina Paz, Cary Hudson, Cedric Whitney, Chak Li, Charles Chen, Charlotte Cole, Chelsea Voss, Chen Ding, Chen Shen, Chengdu Huang, Chris Colby, Chris Hallacy, Chris Koch, Chris Lu, Christina Kaplan, Christina Kim, CJ Minott-Henriques, Cliff Frey, Cody Yu, Coley Czarnecki, Colin Reid, Colin Wei, Cory Decareaux, Cristina Scheau, Cyril Zhang, Cyrus Forbes, Da Tang, Dakota Goldberg, Dan Roberts, Dana Palmie, Daniel Kappler,

Daniel Levine, Daniel Wright, Dave Leo, David Lin, David Robinson, Declan Grabb, Derek Chen, Derek Lim, Derek Salama, Dibya Bhattacharjee, Dimitris Tsipras, Dinghua Li, Dingli Yu, DJ Strouse, Drew Williams, Dylan Hunn, Ed Bayes, Edwin Arbus, Ekin Akyurek, Elaine Ya Le, Elana Widmann, Eli Yani, Elizabeth Proehl, Enis Sert, Enoch Cheung, Eri Schwartz, Eric Han, Eric Jiang, Eric Mitchell, Eric Sigler, Eric Wallace, Erik Ritter, Erin Kavanaugh, Evan Mays, Evgenii Nikishin, Fangyuan Li, Felipe Petroski Such, Filipe de Avila Belbute Peres, Filippo Raso, Florent Bekerma, Foivos Tsim-pourlas, Fotis Chantzis, Francis Song, Francis Zhang, Gaby Raila, Garrett McGrath, Gary Briggs, Gary Yang, Giambattista Parascandolo, Gildas Chabot, Grace Kim, Grace Zhao, Gregory Valiant, Guillaume Leclerc, Hadi Salman, Hanson Wang, Hao Sheng, Haoming Jiang, Haoyu Wang, Haozhun Jin, Harshit Sikchi, Heather Schmidt, Henry Aspegren, Honglin Chen, Huida Qiu, Hunter Lightman, Ian Covert, Ian Kivlichan, Ian Silber, Ian Sohl, Ibrahim Hammoud, Ignasi Clavera, Ikai Lan, Ilge Akkaya, Ilya Kostrikov, Irina Kofman, Isak Etinger, Ishaan Singal, Jackie Hehir, Jacob Huh, Jacqueline Pan, Jake Wilczynski, Jakub Pachocki, James Lee, James Quinn, Jamie Kiros, Janvi Kalra, Jas-myn Samaroo, Jason Wang, Jason Wolfe, Jay Chen, Jay Wang, Jean Harb, Jeffrey Han, Jeffrey Wang, Jennifer Zhao, Jeremy Chen, Jerene Yang, Jerry Tworek, Jesse Chand, Jes-sica Landon, Jessica Liang, Ji Lin, Jiancheng Liu, Jianfeng Wang, Jie Tang, Jihan Yin, Joanne Jang, Joel Morris, Joey Flynn, Johannes Ferstad, Johannes Heidecke, John Fish-bein, John Hallman, Jonah Grant, Jonathan Chien, Jonathan Gordon, Jongsoo Park, Jordan Liss, Jos Kraaijeveld, Joseph Guay, Joseph Mo, Josh Lawson, Josh McGrath, Joshua Vendrow, Joy Jiao, Julian Lee, Julie Steele, Julie Wang, Junhua Mao, Kai Chen, Kai Hayashi, Kai Xiao, Kamyar Salahi, Kan Wu, Karan Sekhri, Karan Sharma, Karan Singhal, Karen Li, Kenny Nguyen, Keren Gu-Lemberg, Kevin King, Kevin Liu, Kevin Stone, Kevin Yu, Kristen Ying, Kristian Georgiev, Kristie Lim, Kushal Tirumala, Kyle Miller, Lama Ahmad, Larry Lv, Laura Clare, Laurance Fauconnet, Lauren Itow, Lauren Yang, Laurentia Romaniuk, Leah Anise, Lee Byron, Leher Pathak, Leon Maksin, Leyan Lo, Leyton Ho, Li Jing, Liang Wu, Liang Xiong, Lien Mamitsuka, Lin Yang, Lindsay McCallum, Lindsey Held, Liz Bourgeois, Logan Engstrom, Lorenz Kuhn, Louis Feувrier, Lu Zhang, Lucas Switzer, Lukas Kondraciuk, Lukasz Kaiser, Manas Joglekar, Mandeep Singh, Mandip Shah, Manuka Stratta, Marcus Williams, Mark Chen, Mark Sun, Marselus Cayton, Martin Li, Marvin Zhang, Marwan Aljubeih, Matt Nichols, Matthew Haines, Max Schwarzer, Mayank Gupta, Meghan Shah, Melody Huang, Meng Dong, Mengqing Wang, Mia Glaese, Micah Carroll, Michael Lampe, Michael Malek, Michael Sharman, Michael Zhang, Michele Wang, Michelle Pokrass, Mihai Florian, Mikhail Pavlov, Miles Wang, Ming Chen, Mingxuan Wang, Minnia Feng, Mo Bavarian, Molly Lin, Moose Abdool, Mostafa Rohaninejad, Nacho Soto, Natalie Staudacher, Natan LaFontaine, Nathan Mar-well, Nelson Liu, Nick Preston, Nick Turley, Nicklas Ansmann, Nicole Blades, Nikil Pancha, Nikita Mikhaylin, Niko Felix, Nikunj Handa, Nishant Rai, Nitish Keskar, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Oona Gleeson, Pamela Mishkin, Patryk Lesiewicz, Paul Baltescu, Pavel Belov, Peter Zhokhov, Philip Pronin, Phillip Guo, Phoebe Thacker, Qi Liu, Qiming Yuan, Qinghua Liu, Rachel Dias, Rachel Puckett, Rahul Arora, Ravi Teja Mullapudi, Raz Gaon, Reah Miyara, Rennie Song, Rishabh Aggarwal, RJ Marsan, Robel Yemiru, Robert Xiong, Rohan Kshirsagar, Rohan Nuttall, Roman Tsiupa, Ronen Eldan, Rose Wang, Roshan James, Roy Ziv, Rui Shu, Ruslan Nigmat-ullin, Saachi Jain, Saam Talaie, Sam Altman, Sam Arnesen, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Sarah Yoo, Savannah Heon, Scott Ethersmith, Sean Grove, Sean Taylor, Sebastien Bubeck, Sever Banesiu, Shaokyi Amdo, Shengjia Zhao, Sherwin Wu, Shibani Santurkar, Shiyu Zhao, Shraman Ray Chaudhuri, Shreyas Krishnaswamy, Shuaiqi, Xia, Shuyang Cheng, Shyamal Anadkat, Simón Posada Fish-man, Simon Tobin, Siyuan Fu, Somay Jain, Song Mei, Sonya Egoian, Spencer Kim, Spug Golden, SQ Mah, Steph Lin, Stephen Imm, Steve Sharpe, Steve Yadlowsky, Sulman Choudhry, Sungwon Eum, Suvansh Sanjeev, Tabarak Khan, Tal Stramer, Tao Wang, Tao Xin, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Degry, Thomas Shadwell, Tianfu Fu, Tianshi Gao, Timur Garipov, Tina Sriskandarajah, Toki Sherbakov, Tomer Kaftan, Tomo Hiratsuka, Tongzhou Wang, Tony Song, Tony Zhao, Troy Peterson, Val Kharitonov, Victoria Chernova, Vineet Kosaraju, Vishal Kuo, Vitchyr Pong, Vivek Verma, Vlad Petrov, Wanning Jiang, Weixing Zhang, Wenda Zhou, Wen-lei Xie, Wenting Zhan, Wes McCabe, Will DePue, Will Ellsworth, Wulfie Bain, Wyatt

Thompson, Xiangning Chen, Xiangyu Qi, Xin Xiang, Xinwei Shi, Yann Dubois, Yaodong Yu, Yara Khakbaz, Yifan Wu, Yilei Qian, Yin Tat Lee, Yinbo Chen, Yizhen Zhang, Yizhong Xiong, Yonglong Tian, Young Cha, Yu Bai, Yu Yang, Yuan Yuan, Yuanzhi Li, Yufeng Zhang, Yuguang Yang, Yujia Jin, Yun Jiang, Yunyun Wang, Yushi Wang, Yutian Liu, Zach Stubenvoll, Zehao Dou, Zheng Wu, and Zhigang Wang. Openai gpt-5 system card, 2025. URL <https://arxiv.org/abs/2601.03267>.

Charlie Snell, Dan Klein, and Ruiqi Zhong. Learning by distilling context, 2022. URL <https://arxiv.org/abs/2209.15189>.

Alessandro Sordani, Xingdi Yuan, Marc-Alexandre Côté, Matheus Pereira, Adam Trischler, Ziang Xiao, Arian Hosseini, Friederike Niedtner, and Nicolas Le Roux. Joint prompt optimization of stacked llms using variational inference, 2023. URL <https://arxiv.org/abs/2306.12509>.

Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Jack Krawczyk, Cosmo Du, Ed Chi, Heng-Tze Cheng, Eric Ni, Purvi Shah, Patrick Kane, Betty Chan, Manaal Faruqui, Aliaksei Severyn, Hanzhao Lin, YaGuang Li, Yong Cheng, Abe Ittycheriah, Mahdis Mahdieh, Mia Chen, Pei Sun, Dustin Tran, Sumit Bagri, Balaji Lakshminarayanan, Jeremiah Liu, Andras Orban, Fabian Gura, Hao Zhou, Xinying Song, Aurelien Boffy, Harish Ganapathy, Steven Zheng, HyunJeong Choe, goston Weisz, Tao Zhu, Yifeng Lu, Siddharth Gopal, Jarrod Kahn, Maciej Kula, Jeff Pitman, Rushin Shah, Emanuel Taropa, Majd Al Merey, Martin Baeuml, Zhifeng Chen, Laurent El Shafey, Yujing Zhang, Olcan Sercinoglu, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anas White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, Alexandre Frechette, Charlotte Smith, Laura Culp, Lev Proleev, Yi Luan, Xi Chen, James Lottes, Nathan Schucher, Federico Lebron, Alban Rrustemi, Natalie Clay, Phil Crone, Tomas Kocisky, Jeffrey Zhao, Bartek Perz, Dian Yu, Heidi Howard, Adam Bloniarz, Jack W. Rae, Han Lu, Laurent Sifre, Marcello Maggioni, Fred Alcober, Dan Garrette, Megan Barnes, Shantanu Thakoor, Jacob Austin, Gabriel Barth-Maron, William Wong, Rishabh Joshi, Rahma Chaabouni, Deeni Fatiha, Arun Ahuja, Gaurav Singh Tomar, Evan Senter, Martin Chadwick, Ilya Kornakov, Nithya Attaluri, Inaki Iturrate, Ruibo Liu, Yunxuan Li, Sarah Cogan, Jeremy Chen, Chao Jia, Chenjie Gu, Qiao Zhang, Jordan Grimstad, Ale Jakse Hartman, Xavier Garcia, Thanumalayan Sankaranarayana Pillai, Jacob Devlin, Michael Laskin, Diego de Las Casas, Dasha Valter, Connie Tao, Lorenzo Blanco, Adri Puigdomnech Badia, David Reitter, Mianna Chen, Jenny Brennan, Clara Rivera, Sergey Brin, Shariq Iqbal, Gabriela Surita, Jane Labanowski, Abhi Rao, Stephanie Winkler, Emilio Parisotto, Yiming Gu, Kate Olszewska, Ravi Addanki, Antoine Miech, Annie Louis, Denis Teplyashin, Geoff Brown, Elliot Catt, Jan Balaguer, Jackie Xiang, Pidong Wang, Zoe Ashwood, Anton Briukhov, Albert Webson, Sanjay Ganapathy, Smit Sanghavi, Ajay Kannan, Ming-Wei Chang, Axel Stjerngren, Josip Djolonga, Yuting Sun, Ankur Bapna, Matthew Aitchison, Pedram Pejman, Henryk Michalewski, Tianhe Yu, Cindy Wang, Juliette Love, Junwhan Ahn, Dawn Bloxwich, Kehang Han, Peter Humphreys, Thibault Sellam, James Bradbury, Varun Godbole, Sina Samangooei, Bogdan Damoc, Alex Kaskasoli, Sbastien M. R. Arnold, Vijay Vasudevan, Shubham Agrawal, Jason Riesa, Dmitry Lepikhin, Richard Tanburn, Srivatsan Srinivasan, Hyeontaek Lim, Sarah Hodgkinson, Pranav Shyam, Johan Ferret, Steven Hand, Ankush Garg, Tom Le Paine, Jian Li, Yujia Li, Minh Giang, Alexander Neitz, Zaheer Abbas, Sarah York, Machel Reid, Elizabeth Cole, Aakanksha Chowdhery, Dipanjan Das, Dominika Rogoziska, Vitaliy Nikolaev,

Pablo Sprechmann, Zachary Nado, Lukas Zilka, Flavien Prost, Luheng He, Marianne Monteiro, Gaurav Mishra, Chris Welty, Josh Newlan, Dawei Jia, Miltiadis Allamanis, Clara Huiyi Hu, Raoul de Liedekerke, Justin Gilmer, Carl Saroufim, Shruti Rijhwani, Shaobo Hou, Disha Shrivastava, Anirudh Baddepudi, Alex Goldin, Adnan Ozturel, Albin Cassirer, Yunhan Xu, Daniel Sohn, Devendra Sachan, Reinald Kim Amplayo, Craig Swanson, Dessie Petrova, Shashi Narayan, Arthur Guez, Siddhartha Brahma, Jessica Landon, Miteyan Patel, Ruizhe Zhao, Kevin Villela, Luyu Wang, Wenhao Jia, Matthew Rahtz, Mai Giménez, Legg Yeung, James Keeling, Petko Georgiev, Diana Mincu, Boxi Wu, Salem Haykal, Rachel Saputro, Kiran Vodrahalli, James Qin, Zeynep Cankara, Abhanshu Sharma, Nick Fernando, Will Hawkins, Behnam Neyshabur, Solomon Kim, Adrian Hutter, Priyanka Agrawal, Alex Castro-Ros, George van den Driessche, Tao Wang, Fan Yang, Shuo yin Chang, Paul Komarek, Ross McIlroy, Mario Lui, Guodong Zhang, Wael Farhan, Michael Sharman, Paul Natsev, Paul Michel, Yamini Bansal, Siyuan Qiao, Kris Cao, Siamak Shakeri, Christina Butterfield, Justin Chung, Paul Kishan Rubenstein, Shivani Agrawal, Arthur Mensch, Kedar Soparkar, Karel Lenc, Timothy Chung, Aedan Pope, Loren Maggiore, Jackie Kay, Priya Jhakra, Shibo Wang, Joshua Maynez, Mary Phuong, Taylor Tobin, Andrea Tacchetti, Maja Trebacz, Kevin Robinson, Yash Katariya, Sebastian Riedel, Paige Bailey, Kefan Xiao, Nimesh Ghelani, Lora Aroyo, Ambrose Slone, Neil Houlsby, Xuehan Xiong, Zhen Yang, Elena Gribovskaya, Jonas Adler, Mateo Wirth, Lisa Lee, Music Li, Thais Kagohara, Jay Pavagadhi, Sophie Bridgers, Anna Bortsova, Sanjay Ghemawat, Zafarali Ahmed, Tianqi Liu, Richard Powell, Vijay Bolina, Mariko Iinuma, Polina Zablotskaia, James Besley, Da-Woon Chung, Timothy Dozat, Ramona Comanescu, Xiance Si, Jeremy Greer, Guolong Su, Martin Polacek, Raphaël Lopez Kaufman, Simon Tokumine, Hexiang Hu, Elena Buchatskaya, Yingjie Miao, Mohamed Elhawaty, Aditya Siddhant, Nenad Tomasev, Jinwei Xing, Christina Greer, Helen Miller, Shereen Ashraf, Aurko Roy, Zizhao Zhang, Ada Ma, Angelos Filos, Milos Besta, Rory Blevins, Ted Klimenko, Chih-Kuan Yeh, Soravit Changpinyo, Jiaqi Mu, Oscar Chang, Mantas Pajarskas, Carrie Muir, Vered Cohen, Charline Le Lan, Krishna Haridasan, Amit Marathe, Steven Hansen, Sholto Douglas, Rajkumar Samuel, Mingqiu Wang, Sophia Austin, Chang Lan, Jiepu Jiang, Justin Chiu, Jaime Alonso Lorenzo, Lars Lowe Sjösund, Sébastien Cevey, Zach Gleicher, Thi Avrahami, Anudhyan Boral, Hansa Srinivasan, Vittorio Selo, Rhys May, Konstantinos Aisopos, Léonard Hussenot, Livio Baldini Soares, Kate Baumli, Michael B. Chang, Adrià Recasens, Ben Caine, Alexander Pritzel, Filip Pavetic, Fabio Pardo, Anita Gergely, Justin Frye, Vinay Ramasesh, Dan Horgan, Kartikeya Badola, Nora Kassner, Subhrajit Roy, Ethan Dyer, Víctor Campos Campos, Alex Tomala, Yunhao Tang, Dalia El Badawy, Elspeth White, Basil Mustafa, Oran Lang, Abhishek Jindal, Sharad Vikram, Zhitao Gong, Sergi Caelles, Ross Hemsley, Gregory Thornton, Fangxiaoyu Feng, Wojciech Stokowiec, Ce Zheng, Phoebe Thacker, Çalar Ünlü, Zhishuai Zhang, Mohammad Saleh, James Svensson, Max Bileschi, Piyush Patil, Ankesh Anand, Roman Ring, Katerina Tsihlias, Arpi Vezer, Marco Selvi, Toby Shevlane, Mikel Rodriguez, Tom Kwiatakowski, Samira Daruki, Keran Rong, Allan Dafoe, Nicholas FitzGerald, Keren Gu-Lemberg, Mina Khan, Lisa Anne Hendricks, Marie Pellat, Vladimir Feinberg, James Cobon-Kerr, Tara Sainath, Maribeth Rauh, Sayed Hadi Hashemi, Richard Ives, Yana Hasson, Eric Noland, Yuan Cao, Nathan Byrd, Le Hou, Qingze Wang, Thibault Sottiaux, Michela Paganini, Jean-Baptiste Lespiau, Alexandre Moufarek, Samer Hassan, Kaushik Shivakumar, Joost van Amersfoort, Amol Mandhane, Pratik Joshi, Anirudh Goyal, Matthew Tung, Andrew Brock, Hannah Sheahan, Vedant Misra, Cheng Li, Nemanja Rakievi, Mostafa Dehghani, Fangyu Liu, Sid Mittal, Junhyuk Oh, Seb Noury, Eren Sezener, Fantine Huot, Matthew Lamm, Nicola De Cao, Charlie Chen, Sidharth Mudgal, Romina Stella, Kevin Brooks, Gautam Vasudevan, Chenxi Liu, Mainak Chain, Nivedita Melinkeri, Aaron Cohen, Venus Wang, Kristie Seymore, Sergey Zubkov, Rahul Goel, Summer Yue, Sai Krishnakumaran, Brian Albert, Nate Hurley, Motoki Sano, Anhad Mohanney, Jonah Joughin, Egor Filonov, Tomasz Kpa, Yomna Eldawy, Jiawern Lim, Rahul Rishi, Shirin Badiezadegan, Taylor Bos, Jerry Chang, Sanil Jain, Sri Gayatri Sundara Padmanabhan, Subha Puttagunta, Kalpesh Krishna, Leslie Baker, Norbert Kalb, Vamsi Bedapudi, Adam Kurzrok, Shuntong Lei, Anthony Yu, Oren Litvin, Xiang Zhou, Zhichun Wu, Sam Sobell, Andrea Siciliano, Alan Papir, Robby Neale, Jonas Bragagnolo, Tej Toor, Tina Chen, Valentin Anklin, Feiran Wang, Richie Feng, Milad Gholami, Kevin Ling, Lijuan Liu, Jules Walter, Hamid Moghaddam, Arun Kishore, Jakub Adamek, Tyler Mer-

cado, Jonathan Mallinson, Siddhinita Wandekar, Stephen Cagle, Eran Ofek, Guillermo Garrido, Clemens Lombriser, Maksim Mukha, Botu Sun, Hafeezul Rahman Mohammad, Josip Matak, Yadi Qian, Vikas Peswani, Pawel Janus, Quan Yuan, Leif Schelin, Oana David, Ankur Garg, Yifan He, Oleksii Duzhyi, Anton Älgmyr, Timothée Lottaz, Qi Li, Vikas Yadav, Luyao Xu, Alex Chinien, Rakesh Shivanna, Aleksandr Chuklin, Josie Li, Carrie Spadine, Travis Wolfe, Kareem Mohamed, Subhabrata Das, Zihang Dai, Kyle He, Daniel von Dincklage, Shyam Upadhyay, Akanksha Maurya, Luyan Chi, Sebastian Krause, Khalid Salama, Pam G Rabinovitch, Pavan Kumar Reddy M, Aarush Selvan, Mikhail Dektiarev, Golnaz Ghiasi, Erdem Guven, Himanshu Gupta, Boyi Liu, Deepak Sharma, Idan Heimlich Shtacher, Shachi Paul, Oscar Akerlund, François-Xavier Aubet, Terry Huang, Chen Zhu, Eric Zhu, Elico Teixeira, Matthew Fritze, Francesco Bertolini, Liana-Eleonora Marinescu, Martin Bölle, Dominik Paulus, Khyatti Gupta, Tejasi Latkar, Max Chang, Jason Sanders, Roopa Wilson, Xuewei Wu, Yi-Xuan Tan, Lam Nguyen Thiet, Tulsee Doshi, Sid Lall, Swaroop Mishra, Wanming Chen, Thang Luong, Seth Benjamin, Jasmine Lee, Ewa Andrejczuk, Dominik Rabiej, Vipul Ranjan, Krzysztof Styrz, Pengcheng Yin, Jon Simon, Malcolm Rose Harriott, Mudit Bansal, Alexei Robsky, Geoff Bacon, David Greene, Daniil Mirylenka, Chen Zhou, Obaid Sarvana, Abhimanyu Goyal, Samuel Andermatt, Patrick Siegler, Ben Horn, Assaf Israel, Francesco Pongetti, Chih-Wei "Louis" Chen, Marco Selvatici, Pedro Silva, Kathie Wang, Jackson Tolins, Kelvin Guu, Roey Yogev, Xiaochen Cai, Alessandro Agostini, Maulik Shah, Hung Nguyen, Noah Ó Donnaile, Sébastien Pereira, Linda Friso, Adam Stambler, Adam Kurzrok, Chenkai Kuang, Yan Romanikhin, Mark Geller, ZJ Yan, Kane Jang, Cheng-Chun Lee, Wojciech Fica, Eric Malmi, Qijun Tan, Dan Banica, Daniel Balle, Ryan Pham, Yanping Huang, Diana Avram, Hongzhi Shi, Jasjot Singh, Chris Hidey, Niharika Ahuja, Pranab Saxena, Dan Dooley, Srividya Pranavi Potharaju, Eileen O'Neill, Anand Gokulchandran, Ryan Foley, Kai Zhao, Mike Dusenberry, Yuan Liu, Pulkit Mehta, Ragha Kotikalapudi, Chalence Safranek-Shrader, Andrew Goodman, Joshua Kessinger, Eran Globen, Prateek Kolhar, Chris Gorgolewski, Ali Ibrahim, Yang Song, Ali Eichenbaum, Thomas Brovelli, Sahitya Potluri, Preethi Lahoti, Cip Baetu, Ali Ghorbani, Charles Chen, Andy Crawford, Shalini Pal, Mukund Sridhar, Petru Gurita, Asier Mujika, Igor Petrovski, Pierre-Louis Cedoz, Chenmei Li, Shiyuan Chen, Niccolò Dal Santo, Siddharth Goyal, Jitesh Punjabi, Karthik Kappaganthu, Chester Kwak, Pallavi LV, Sarmishta Velury, Himadri Choudhury, Jamie Hall, Premal Shah, Ricardo Figueira, Matt Thomas, Minjie Lu, Ting Zhou, Chintu Kumar, Thomas Jurdi, Sharat Chikkerur, Yenai Ma, Adams Yu, Soo Kwak, Victor Åhdel, Sujeevan Rajayogam, Travis Choma, Fei Liu, Aditya Barua, Colin Ji, Ji Ho Park, Vincent Hellendoorn, Alex Bailey, Taylan Bilal, Huanjie Zhou, Mehrdad Khatir, Charles Sutton, Wojciech Rządowski, Fiona Macintosh, Roopali Vij, Konstantin Shagin, Paul Medina, Chen Liang, Jinjing Zhou, Pararth Shah, Yingying Bi, Attila Dankovics, Shipra Banga, Sabine Lehmann, Marissa Bredesen, Zifan Lin, John Eric Hoffmann, Jonathan Lai, Raynald Chung, Kai Yang, Nihal Balani, Arthur Brainskas, Andrei Sozanschi, Matthew Hayes, Héctor Fernández Alcalde, Peter Makarov, Will Chen, Antonio Stella, Liselotte Snijders, Michael Mandl, Ante Kärman, Pawe Nowak, Xinyi Wu, Alex Dyck, Krishnan Vaidyanathan, Raghavender R, Jessica Mallet, Mitch Rudominer, Eric Johnston, Sushil Mittal, Akhil Udathu, Janara Christensen, Vishal Verma, Zach Irving, Andreas Santucci, Gamaleldin Elsayed, Elnaz Davoodi, Marin Georgiev, Ian Tenney, Nan Hua, Geoffrey Cideron, Edouard Leurent, Mahmoud Alnahlawi, Ionut Georgescu, Nan Wei, Ivy Zheng, Dylan Scandinaro, Heinrich Jiang, Jasper Snoek, Mukund Sundararajan, Xuezhi Wang, Zack Ontiveros, Itay Karo, Jeremy Cole, Vinu Rajashekhar, Lara Tumeh, Eyal Ben-David, Rishub Jain, Jonathan Uesato, Romina Datta, Oskar Bunyan, Shimu Wu, John Zhang, Piotr Stanczyk, Ye Zhang, David Steiner, Subhajt Naskar, Michael Azzam, Matthew Johnson, Adam Paszke, Chung-Cheng Chiu, Jaume Sanchez Elias, Afroz Mohiuddin, Faizan Muhammad, Jin Miao, Andrew Lee, Nino Vieillard, Jane Park, Jiageng Zhang, Jeff Stanway, Drew Garmon, Abhijit Karmarkar, Zhe Dong, Jong Lee, Aviral Kumar, Luowei Zhou, Jonathan Evens, William Isaac, Geoffrey Irving, Edward Loper, Michael Fink, Isha Arkatkar, Nanxin Chen, Izhak Shafran, Ivan Ptrychenko, Zhe Chen, Johnson Jia, Anselm Levskaya, Zhenkai Zhu, Peter Grabowski, Yu Mao, Alberto Magni, Kaisheng Yao, Javier Snaider, Norman Casagrande, Evan Palmer, Paul Suganthan, Alfonso Castaño, Irene Giannoumis, Wooyeol Kim, Mikoaj Rybiski, Ashwin Sreevatsa, Jennifer Prendki, David Soergel, Adrian Goedeckemeyer, Willi Gierke, Mohsen Jafari, Meenu Gaba, Jeremy Wies-

ner, Diana Gage Wright, Yawen Wei, Harsha Vashisht, Yana Kulizhskaya, Jay Hoover, Maigo Le, Lu Li, Chimezie Iwuanyanwu, Lu Liu, Kevin Ramirez, Andrey Khorlin, Albert Cui, Tian LIN, Marcus Wu, Ricardo Aguilar, Keith Pallo, Abhishek Chakladar, Ginger Perng, Elena Allica Abellan, Mingyang Zhang, Ishita Dasgupta, Nate Kushman, Ivo Penchev, Alena Repina, Xihui Wu, Tom van der Weide, Priya Ponnappalli, Caroline Kaplan, Jiri Simsa, Shuangfeng Li, Olivier Dousse, Fan Yang, Jeff Piper, Nathan Ie, Rama Pasumarthi, Nathan Lintz, Anitha Vijayakumar, Daniel Andor, Pedro Valenzuela, Minnie Lui, Cosmin Paduraru, Daiyi Peng, Katherine Lee, Shuyuan Zhang, Somer Greene, Duc Dung Nguyen, Paula Kurylowicz, Cassidy Hardin, Lucas Dixon, Lili Janzer, Kiam Choo, Ziqiang Feng, Biao Zhang, Achintya Singhal, Dayou Du, Dan McKinnon, Natasha Antropova, Tolga Bolukbasi, Orgad Keller, David Reid, Daniel Finchelstein, Maria Abi Raad, Remi Crocker, Peter Hawkins, Robert Dadashi, Colin Gaffney, Ken Franko, Anna Bulanova, Rémi Leblond, Shirley Chung, Harry Askham, Luis C. Cobo, Kelvin Xu, Felix Fischer, Jun Xu, Christina Sorokin, Chris Alberti, Chu-Cheng Lin, Colin Evans, Alek Dimitriev, Hannah Forbes, Dylan Banarse, Zora Tung, Mark Omernick, Colton Bishop, Rachel Sterneck, Rohan Jain, Jiawei Xia, Ehsan Amid, Francesco Piccinno, Xingyu Wang, Praseem Banzal, Daniel J. Mankowitz, Alex Polozov, Victoria Krakovna, Sasha Brown, MohammadHossein Bateni, Dennis Duan, Vlad Firoiu, Meghana Thotakuri, Tom Natan, Matthieu Geist, Ser tan Girgin, Hui Li, Jiayu Ye, Ofir Roval, Reiko Tojo, Michael Kwong, James Lee-Thorp, Christopher Yew, Danila Sinopalnikov, Sabela Ramos, John Mellor, Abhishek Sharma, Kathy Wu, David Miller, Nicolas Sonnerat, Denis Vnukov, Rory Greig, Jennifer Beattie, Emily Caveness, Libin Bai, Julian Eisenschlos, Alex Korchemniy, Tomy Tsai, Mimi Jasarevic, Weize Kong, Phuong Dao, Zeyu Zheng, Frederick Liu, Fan Yang, Rui Zhu, Tian Huey Teh, Jason Sanmiya, Evgeny Gladchenko, Nejc Trdin, Daniel Toyama, Evan Rosen, Sasan Tavakkol, Linting Xue, Chen Elkind, Oliver Woodman, John Carpenter, George Papamakarios, Rupert Kemp, Sushant Kafle, Tanya Grunina, Rishika Sinha, Alice Talbert, Diane Wu, Denese Owusu-Afriyie, Cosmo Du, Chloe Thornton, Jordi Pont-Tuset, Pradyumna Narayana, Jing Li, Saaber Fatehi, John Wieting, Omar Ajmeri, Benigno Uria, Yeongil Ko, Laura Knight, Amélie Héliou, Ning Niu, Shane Gu, Chenxi Pang, Yeqing Li, Nir Levine, Ariel Stolovich, Rebeca Santamaria-Fernandez, Sonam Goenka, Wenny Yustalim, Robin Strudel, Ali Elqursh, Charlie Deck, Hyo Lee, Zonglin Li, Kyle Levin, Raphael Hoffmann, Dan Holtmann-Rice, Olivier Bachem, Sho Arora, Christy Koh, Soheil Hassas Yeganeh, Siim Pöder, Mukarram Tariq, Yanhua Sun, Lucian Ionita, Mojtaba Seyedhosseini, Pouya Tafti, Zhiyu Liu, Anmol Gulati, Jasmine Liu, Xinyu Ye, Bart Chrzaszcz, Lily Wang, Nikhil Sethi, Tianrun Li, Ben Brown, Shreya Singh, Wei Fan, Aaron Parisi, Joe Stanton, Vinod Koverkathu, Christopher A. Choquette-Choo, Yunjie Li, TJ Lu, Abe Ittycheriah, Prakash Shroff, Mani Varadarajan, Sanaz Bahargam, Rob Willoughby, David Gaddy, Guillaume Desjardins, Marco Cornero, Brona Robenek, Bhavishya Mittal, Ben Albrecht, Ashish Shenoy, Fedor Moiseev, Henrik Jacobsson, Alireza Ghaffarkhah, Morgane Rivière, Alanna Walton, Clément Crepy, Alicia Parrish, Zongwei Zhou, Clement Farabet, Carey Radebaugh, Praveen Srinivasan, Claudia van der Salm, Andreas Fijdeland, Salvatore Scellato, Eri Latorre-Chimoto, Hanna Klimczak-Pluciska, David Bridson, Dario de Cesare, Tom Hudson, Piermaria Mendolicchio, Lexi Walker, Alex Morris, Matthew Mauger, Alexey Guseynov, Alison Reid, Seth Odoom, Lucia Lohrer, Victor Cotruta, Madhavi Yenugula, Dominik Grewe, Anastasia Petrushkina, Tom Duerig, Antonio Sanchez, Steve Yadlowsky, Amy Shen, Amir Globerson, Lynette Webb, Sahil Dua, Dong Li, Surya Bhupatiraju, Dan Hurt, Haroon Qureshi, Ananth Agarwal, Tomer Shani, Matan Eyal, Anuj Khare, Shreyas Rammohan Belle, Lei Wang, Chetan Tekur, Mihir Sanjay Kale, Jinliang Wei, Ruoxin Sang, Brennan Saeta, Tyler Liechty, Yi Sun, Yao Zhao, Stephan Lee, Pandu Nayak, Doug Fritz, Manish Reddy Vuyyuru, John Aslanides, Nidhi Vyas, Martin Wicke, Xiao Ma, Evgenii Eltyshev, Nina Martin, Hardie Cate, James Manyika, Keyvan Amiri, Yelin Kim, Xi Xiong, Kai Kang, Florian Luisier, Nilesch Tripuraneni, David Madras, Mandy Guo, Austin Waters, Oliver Wang, Joshua Ainslie, Jason Baldridge, Han Zhang, Garima Pruthi, Jakob Bauer, Feng Yang, Riham Mansour, Jason Gelman, Yang Xu, George Polovets, Ji Liu, Honglong Cai, Warren Chen, XiangHai Sheng, Emily Xue, Sherjil Ozair, Christof Angermueller, Xiaowei Li, Anoop Sinha, Weiren Wang, Julia Wiesinger, Emmanouil Koukoumidis, Yuan Tian, Anand Iyer, Madhu Gurusurthy, Mark Goldenson, Parashar Shah, MK Blake, Hongkun Yu, Anthony Urbanowicz, Jennimaria Palomaki, Chrisantha Fernando, Ken Durden, Harsh Mehta,

- Nikola Momchev, Elahe Rahimtoroghi, Maria Georgaki, Amit Raul, Sebastian Ruder, Morgan Redshaw, Jinhyuk Lee, Denny Zhou, Komal Jalan, Dinghua Li, Blake Hechtman, Parker Schuh, Milad Nasr, Kieran Milan, Vladimir Mikulik, Juliana Franco, Tim Green, Nam Nguyen, Joe Kelley, Aroma Mahendru, Andrea Hu, Joshua Howland, Ben Vargas, Jeffrey Hui, Kshitij Bansal, Vikram Rao, Rakesh Ghiya, Emma Wang, Ke Ye, Jean Michel Sarr, Melanie Moranski Preston, Madeleine Elish, Steve Li, Aakash Kaku, Jigar Gupta, Ice Pasupat, Da-Cheng Juan, Milan Someswar, Tejvi M., Xinyun Chen, Aida Amini, Alex Fabrikant, Eric Chu, Xuanyi Dong, Amruta Muthal, Senaka Buthpitiya, Sarthak Jauhari, Nan Hua, Urvashi Khandelwal, Ayal Hitron, Jie Ren, Larissa Rinaldi, Shahar Drath, Avigail Dabush, Nan-Jiang Jiang, Harshal Godhia, Uli Sachs, Anthony Chen, Yicheng Fan, Hagai Taitelbaum, Hila Noga, Zhuyun Dai, James Wang, Chen Liang, Jenny Hamer, Chun-Sung Ferng, Chenel Elkind, Aviel Atias, Paulina Lee, Vít Listík, Mathias Carlen, Jan van de Kerkhof, Marcin Pikus, Krunoslav Zaher, Paul Müller, Sasha Zykova, Richard Stefanec, Vitaly Gatsko, Christoph Hirsenschall, Ashwin Sethi, Xingyu Federico Xu, Chetan Ahuja, Beth Tsai, Anca Stefanoiu, Bo Feng, Keshav Dhandhanania, Manish Katyal, Akshay Gupta, Atharva Parulekar, Divya Pitta, Jing Zhao, Vivaan Bhatia, Yashodha Bhavnani, Omar Alhadlaq, Xiaolin Li, Peter Danenberg, Dennis Tu, Alex Pine, Vera Filippova, Abhipso Ghosh, Ben Limonchik, Bhargava Urala, Chaitanya Krishna Lanka, Derik Clive, Yi Sun, Edward Li, Hao Wu, Kevin Hongtongsak, Ianna Li, Kalind Thakkar, Kuanysh Omarov, Kushal Majmudar, Michael Alverson, Michael Kucharski, Mohak Patel, Mudit Jain, Maksim Zabelin, Paolo Pelagatti, Rohan Kohli, Saurabh Kumar, Joseph Kim, Swetha Sankar, Vineet Shah, Lakshmi Ramachandruni, Xiangkai Zeng, Ben Bariach, Laura Weidinger, Tu Vu, Alek Andreev, Antoine He, Kevin Hui, Sheleem Kashem, Amar Subramanya, Sissie Hsiao, Demis Hassa. Gemini: A family of highly capable multimodal models, 2025. URL <https://arxiv.org/abs/2312.11805>.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. MuSiQue: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554, 2022.
- Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22(5–6):544–557, 2009. doi: 10.1016/j.neunet.2009.06.042. URL <https://www.sciencedirect.com/science/article/abs/pii/S0893608009001130>.
- Dheeraj Vattikonda, Santhoshi Ravichandran, Emiliano Penaloza, Hadi Nekoei, Megh Thakkar, Thibault Le Sellier de Chezelles, Nicolas Gontier, Miguel Muñoz-Mármol, Sahar Omid Shayegan, Stefania Raimondo, Xue Liu, Alexandre Drouin, Laurent Charlin, Alexandre Piché, Alexandre Lacoste, and Massimo Caccia. How to train your llm web agent: A statistical diagnosis, 2025. URL <https://arxiv.org/abs/2507.04103>.
- Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and Yu Su. Travelplanner: A benchmark for real-world planning with language agents, 2024. URL <https://arxiv.org/abs/2402.01622>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- 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*, pp. 2369–2380, 2018.

- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ -bench: A benchmark for tool-agent-user interaction in real-world domains, 2024. URL <https://arxiv.org/abs/2406.12045>.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-of-thought reasoning in llms, 2025. URL <https://arxiv.org/abs/2502.03373>.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025. URL <https://arxiv.org/abs/2503.14476>.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. Star: Bootstrapping reasoning with reasoning, 2022. URL <https://arxiv.org/abs/2203.14465>.
- Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. Quiet-star: Language models can teach themselves to think before speaking, 2024. URL <https://arxiv.org/abs/2403.09629>.
- 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, 2026. URL <https://arxiv.org/abs/2601.18734>.
- Xiangxin Zhou, Zichen Liu, Haonan Wang, Chao Du, Min Lin, Chongxuan Li, Liang Wang, and Tianyu Pang. Variational reasoning for language models, 2025. URL <https://arxiv.org/abs/2509.22637>.

A RELATED WORK

Latent Reasoning and Variational Perspectives. Recent work increasingly frames LM reasoning as a latent-variable inference problem (Hu et al., 2024; Sordoni et al., 2023; Luo et al., 2025; Li et al., 2025). Within this framework, the most popular approach is STaR (Li et al., 2025), which uses a hint when the model is unable to correctly answer the question and then uses SFT to fit onto the generated reasoning trace. This is similar to student-only training in π -Distill ($\alpha = 0$), where rather than a KL-regularized off-policy RL objective, SFT is used and the teacher is not trained. Most similar to π -Distill is the work by (Zhou et al., 2025), which proposes variational reasoners. This approach can be seen as a modified version of STaR, where both teacher and student are being trained iteratively. The main difference between variational reasoning and our work is they assume access to oracle answers and perform an iterative version of variational EM using separate parameters for the teacher and student. In comparison, we simplify the training objective by allowing the teacher and student to share parameters and do not assume access to ground truth answers. We compare against a similar setup with minor modifications due to not having access to ground truth solutions in App. C.2.

Self-Bootstrapping, Privileged Signals, and Guided Exploration. Complementary work focuses on using privileged or auxiliary signals to enable learning in hard regimes. (Chen et al., 2025) injects self-generated high-level hints into online RL to overcome zero-reward exploration barriers, while (Qu et al., 2026) uses privileged oracle solutions as structured on-policy exploration signals for hard reasoning tasks. Both these algorithms can be viewed as training the teacher to use PI with implicit transfer via parameter sharing. Additionally, other lines of work explore distilling certain skills/behaviors in the model via contextualized sampling (Didolkar et al., 2025; 2024). We believe our proposed methods could be used to distill desired skills/behaviors into model weights in a more effective manner compared to traditional SFT.

Context Distillation. Both our proposed methods can be seen as a form of context distillation. This growing line of work studies how to train models to internalize context-dependent reasoning so that the benefits of rich contexts (e.g., instructions, demonstrations, auxiliary computations) are distilled into the model’s weights (Caccia et al., 2025; Snell et al., 2022; Huang et al., 2022; Cao et al., 2025). Other works analyze this problem by performing SFT answers they give when having access to context Snell et al. (2022). Moreover, other works also try to encapsulate such knowledge into adapters that can be plugged and played at test time (Caccia et al., 2025).

B LIMITATIONS & FUTURE WORK

While we show that both π -Distill and OPSD are effective algorithms when CoT is not available, all our PI is transformed from frontier model traces. Exploring how these algorithms perform without access to a frontier model or ground truth answers is an interesting direction. Additionally, our experiments are limited to models ≤ 8 B parameters. Further scaling of both π -Distill and OPSD can shine light on additional important factors to enable these algorithms. Moreover, our analysis in Section 7 is limited to observational studies, where we do not systematically control for all variables, rather observe properties of different configurations. A study that systematically controls for these factors could help in mitigating failure cases and obtaining the most useful PI.

C CONNECTIONS TO VARIATIONAL EM

In this Appendix, we derive $J_{\pi\text{-Distill}}$ using a variational perspective. Specifically we show how $J_{\pi\text{-Distill}}$ can be seen as a joint EM algorithm where E and M are done simultaneously.

Traditional variational-EM aims to fit a parameterized distribution π^S onto a target posterior π^* . An assumption this framework makes is that using π^S to approximate π^* may

be suboptimal, rather leveraging an approximate posterior $\pi_\theta^T(\mathbf{o} | s, \mathbf{I})$ conditioned on additional information can make the approximation easier. This is true when π^T can sample from high-reward modes that π^S cannot, therefore making optimization via RL impossible for π^S . We visualize the procedure in Figure 3

C.1 VARIATIONAL EM

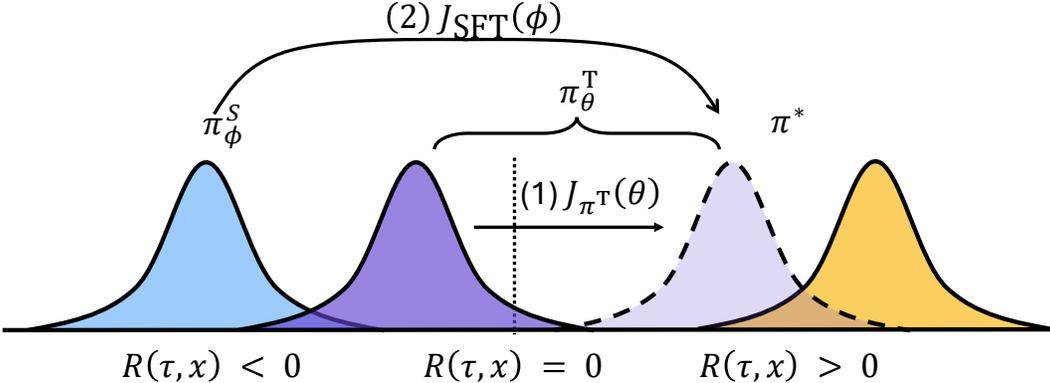


Figure 3: **Illustration of the Variational EM procedure for policy optimization.** The initial student policy, π_ϕ^S , lacks support over trajectories with positive rewards ($R(\mathbf{o}, \mathbf{s}) > 0$), preventing direct improvement. To address this: (1) the teacher policy π_θ^T is optimized via $J_{\text{Teacher}}(\theta)$ to sample successful traces and approximate the optimal policy π^* ; (2) the student policy is then updated via $J_{\text{SFT}}(\theta)$ to distill the knowledge from the teacher. While this two-step procedure is principled, it is computationally inefficient due to the requirement of maintaining dual parameter sets and distinct training phases. In contrast, π -Distill simplifies this pipeline into a single-phase process, providing superior performance with reduced complexity.

E-step J_{Teacher} . We first define the target posterior we want to fit, π^* , as a reward-tilted posterior relative to the reference policy π_{ref} . For a given state \mathbf{s} , we define:

$$\pi^*(\mathbf{o} | \mathbf{s}) = \frac{\pi_{\text{ref}}(\mathbf{o} | \mathbf{s}) \exp(R(\mathbf{o}, \mathbf{s}))}{Z} \tag{6}$$

where $Z = \sum_{\mathbf{o}'} \pi_{\text{ref}}(\mathbf{o}' | \mathbf{s}_0) \exp(R(\mathbf{o}', \mathbf{s}))$ is the partition function and $\mathbf{o} = (\mathbf{z}, \mathbf{a})$ consists of internal reasoning tokens \mathbf{z} and action tokens \mathbf{a} .

The partition function makes this distribution intractable, but it can be approximated using a variational posterior $\pi_\theta^T(\mathbf{o} | s, \mathbf{I})$ conditioned on privileged information \mathbf{I} . We optimize parameters θ by minimizing the reverse KL between the variational distribution and the target policy:

$$\begin{aligned} J_{\text{Teacher}}(\theta) &= -D_{\text{KL}}(\pi_\theta^T(\mathbf{o} | s, \mathbf{I}) \| \pi^*(\cdot | \mathbf{s})) \\ &= -\mathbb{E}_{\substack{\mathbf{o} \sim \pi_\theta^T(\mathbf{o} | s, \mathbf{I}) \\ \mathbf{s} \sim P}} \left[\log \frac{\pi_\theta^T(\mathbf{o} | s, \mathbf{I})}{\pi^*(\mathbf{o} | \mathbf{s})} \right] \\ &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_\theta^T(\mathbf{o} | s, \mathbf{I}) \\ \mathbf{s} \sim P}} \left[\log \pi^*(\mathbf{o} | \mathbf{s}) - \log \pi_\theta^T(\mathbf{o} | s, \mathbf{I}) \right] \end{aligned} \tag{7}$$

Substituting the definition of π^* into the equation:

$$\begin{aligned}
 J_{\text{Teacher}}(\theta) &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} \left[\log \left(\frac{\pi_{\text{ref}}(\mathbf{o} | \mathbf{s}) \exp(R(\mathbf{o}, \mathbf{s}))}{Z} \right) - \log \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \right] \\
 &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} [R(\mathbf{o}, \mathbf{s}) + \log \pi_{\text{ref}}(\mathbf{o} | \mathbf{s}) - \log \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) - \log Z] \\
 &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} [R(\mathbf{o}, \mathbf{s}) - (\log \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) - \log \pi_{\text{ref}}(\mathbf{o} | \mathbf{s}))] - \log Z \\
 &\propto \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} [R(\mathbf{o}, \mathbf{s})] - D_{\text{KL}}(\pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \| \pi_{\text{ref}}(\cdot | \mathbf{s}))
 \end{aligned} \tag{8}$$

Since $\log Z$ depends only on the context \mathbf{s} , it is constant with respect to θ and can be omitted.

M-step J_{SFT} . If π^* were a tractable distribution, we could directly fit the target policy by minimizing the forward KL divergence, $D_{\text{KL}}(\pi^* \| \pi_{\theta}^{\text{S}}$), corresponding to standard SFT. Since π^* is intractable, we instead substitute it with our learned approximation π_{θ}^{T} , yielding the following objective:

$$J_{\text{SFT}}(\theta) = \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{T}}(\mathbf{o} | \mathbf{s}, \mathbf{I}) \\ \mathbf{s} \sim P}} [\log \pi_{\theta}^{\text{S}}(\mathbf{o} | \mathbf{s})]. \tag{9}$$

Traditionally, one would either first fit π_{θ}^{T} to convergence and use it to fit π_{θ}^{S} , or alternate between training the teacher and student. Regardless, this process requires two separate models and can be difficult to optimize. Our implementation is a modified version of Zhou et al. (2025), where the modifications are specific to our multi-turn setting where no ground truth is provided.

C.2 FAILURES OF VARIATIONAL EM.

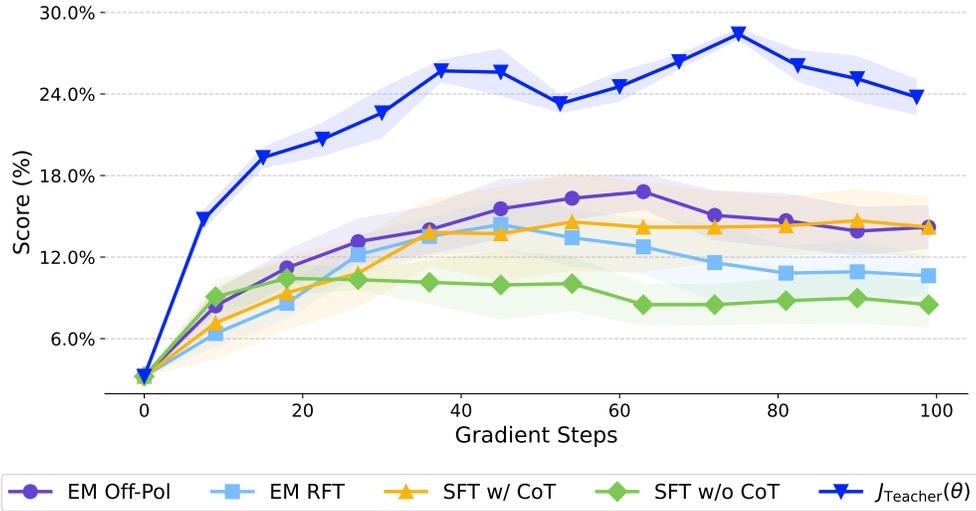


Figure 4: **Performance of Sequential EM variants on τ -Bench Retail (Qwen3-8B).** We compare these methods against SFT baselines. While standard EM RFT (light blue) fails to match SFT w/ CoT (yellow), replacing RFT with off-policy RL (EM Off-Pol, purple) successfully closes allows for EM to outperform SFT w CoT. The strongest performance comes from $J_{\text{Teacher}}(\theta)$ (dark blue), demonstrating that parameter sharing between the teacher and student yields the most effective transfer. Shaded regions indicate standard error across 3 seeds.

Our initial efforts focused on making variational EM work, similar in spirit to Zhou et al. (2025). We first document initial ideas that were unsuccessful, then we discuss the best version of EM we could find being a sequential implementation. Finally we discuss how these experiments ultimately motivated the design decisions for π -Distill.

Initial Experiments First, we find that setting the reference policy π_{ref} as a copy of the current student π_{θ}^{S} is crucial for learning. In contrast, using the fixed base model π_{base} as a prior significantly hinders performance (see App. H.2).

Building on this insight, we initially explored an offline variational EM method similar to Zhou et al. (2025). Here, the goal is learning a CoT sampler as the variational posterior conditioned on expert traces. We do this with the following objective:

$$J_{\text{Offline-EM}}(\theta) = \mathbb{E}_{\mathbf{s} \sim \mathcal{D}} \left[\log \sum_{\mathbf{z}} \pi_{\theta}^{\text{S}}(\mathbf{a}, \mathbf{z} \mid \mathbf{s}) \exp(R(\mathbf{o}, \mathbf{s})) \right] \\ \propto \mathbb{E}_{\mathbf{a}, \mathbf{s} \sim \mathcal{D}, \mathbf{z} \sim \pi_{\theta}^{\text{T}}} [R(\mathbf{o}, \mathbf{s}) + \log \pi_{\theta}^{\text{S}}(\mathbf{a} \mid \mathbf{z}, \mathbf{s})] - D_{\text{KL}}(\pi_{\theta}^{\text{T}}(\mathbf{z} \mid \mathbf{s}) \parallel \pi_{\theta}^{\text{S}}(\mathbf{z} \mid \mathbf{s})) \quad (10)$$

notice here how the states are fixed and given in the offline dataset \mathcal{D} . Furthermore, this setup necessitates substantial hyperparameter tuning. Beyond the standard penalty β , the algorithm requires balancing the ratio of teacher-to-student updates, managing separate optimizer states, and calibrating the dataset size sampled from the teacher. Ultimately, we find this approach hard to implement and fails in our multi-turn setting. We attribute the success of Zhou et al. (2025) to their access to oracle answers; in our context, we find that fitting a variational posterior to suboptimal frontier model trajectories can even degrade performance.

Our following efforts focused on training the teacher online using τ . We try this with alternating loops between student and teacher training, but find this to be ineffective. As we found it hard to avoid the student and teacher collapsing onto each other using this setup. Ultimately, we found a sequential version of EM worked best, though it still requires specific adjustments to achieve optimal performance.

Sequential EM. We turn to Sequential EM after ruling out offline and alternating approaches. However, we find standard implementations remain unreliable for complex agentic tasks. Here we analyze the limitations of standard EM and how they ultimately motivated π -Distill.

For this analysis we use τ -Bench Retail with `Qwen3-8B` using self-generated hints to train the teacher as this configuration performs best (see § 7). We compare sequential EM, where we first train the teacher and then perform SFT on successful trajectories (Rejection Fine-tuning or RFT), against SFT with and without CoT. We note that here it would be sufficient to outperform SFT without CoT as we assume no access to internal CoT of the frontier model, thus in practice sequential EM could replace the traditional SFT phase before RL.

We report results in Figure 4. First, we observe that standard Sequential EM outperforms naive SFT but lags behind training with CoT. We attribute the performance gap to the lack of negative feedback on failed trajectories. It only mimics the teacher. To address this, we replaced the RFT step with clipped off-policy RL. We find this change enables sequential EM to outperform SFT with CoT.

Most surprisingly, in these experiments, we find that simply optimizing $J_{\text{Teacher}}(\theta)$ substantially improves the student π_{θ}^{S} when both teacher and student share parameters. This drastically outperforms all other baselines.

To summarize, our two main findings are that parameter sharing enables substantial transfer from teacher to student and the importance of leveraging negative traces via off-policy RL when training the student. We use these two findings to instantiate π -Distill while also simplifying the process from two training steps to a single one.

D ON-POLICY SELF-DISTILLATION J_{OPSD}

Contrary to the prior objectives, OPSD rather than trying to approximate π^* can be framed as the reverse-KL between π and the target conditioned policy π^{T*} , being defined as:

$$\pi^*(\mathbf{o} \mid \mathbf{s}, \mathbf{I}) = \frac{\pi_{\text{ref}}(\mathbf{o} \mid \mathbf{s}, \mathbf{I}) \exp(\frac{R(\mathbf{o}, \mathbf{s})}{\beta})}{Z^h} \quad (11)$$

where $Z^h = \sum_{\mathbf{o}'} \pi_{\text{ref}}^T(\mathbf{o}' \mid \mathbf{s}, \mathbf{I}) \exp(R(\mathbf{o}', \mathbf{s}))$ is the partition function.

One can then approximate this distribution directly via reverse KL, giving the following objective:

To do this, we minimize the reverse KL between the student policy π_{θ}^{S} and the optimal privileged target distribution π^* :

$$\begin{aligned} D_{\text{KL}}(\pi_{\theta}^{\text{S}} \parallel \pi^*) &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} \left[\log \frac{\pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s})}{\pi^*(\mathbf{o} \mid \mathbf{s}, \mathbf{I})} \right] \\ &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} \left[\log \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) - \log \left(\frac{\pi_{\text{ref}}(\mathbf{o} \mid \mathbf{s}, \mathbf{I}) \exp(R(\mathbf{o}, \mathbf{s})/\beta)}{Z^h} \right) \right] \\ &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} \left[\log \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) - \log \pi_{\text{ref}}(\mathbf{o} \mid \mathbf{s}, \mathbf{I}) - \frac{R(\mathbf{o}, \mathbf{s})}{\beta} + \log Z^h \right] \\ &= \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} \left[-\frac{R(\mathbf{o}, \mathbf{s})}{\beta} + \log \frac{\pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s})}{\pi_{\text{ref}}(\mathbf{o} \mid \mathbf{s}, \mathbf{I})} \right] + \log Z^h \quad (12) \end{aligned}$$

$$= -\frac{1}{\beta} \mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} [R(\mathbf{o}, \mathbf{s})] + D_{\text{KL}}(\pi_{\theta}^{\text{S}}(\cdot \mid \mathbf{s}) \parallel \pi_{\text{ref}}(\cdot \mid \mathbf{s}, \mathbf{I})) + \log Z^h. \quad (13)$$

$$\propto -\mathbb{E}_{\substack{\mathbf{o} \sim \pi_{\theta}^{\text{S}}(\mathbf{o} \mid \mathbf{s}) \\ \mathbf{s} \sim P}} [R(\mathbf{o}, \mathbf{s})] + \beta D_{\text{KL}}(\pi_{\theta}^{\text{S}}(\cdot \mid \mathbf{s}) \parallel \pi_{\text{ref}}(\cdot \mid \mathbf{s}, \mathbf{I})) \quad (14)$$

To align with π -Distill and the according to the results in App. H.2, we set $\pi_{\text{ref}} = \pi_{\theta}^{\text{T}}(\mathbf{o} \mid \mathbf{s}, \mathbf{I})$

Concurrent On-Policy Distillation Methods. Concurrent work proposes closely related on-policy self-distillation frameworks. Recent methods train student models on their own trajectories while using privileged or conditioned variants of the same model as teachers, enabling on-policy transfer without off-policy distribution shift (Zhao et al., 2026; Shenfeld et al., 2026). These approaches are conceptually aligned with our OCPD formulation, but target supervised reasoning or continual learning settings, whereas our focus is agentic decision-making with training-time privileged information.

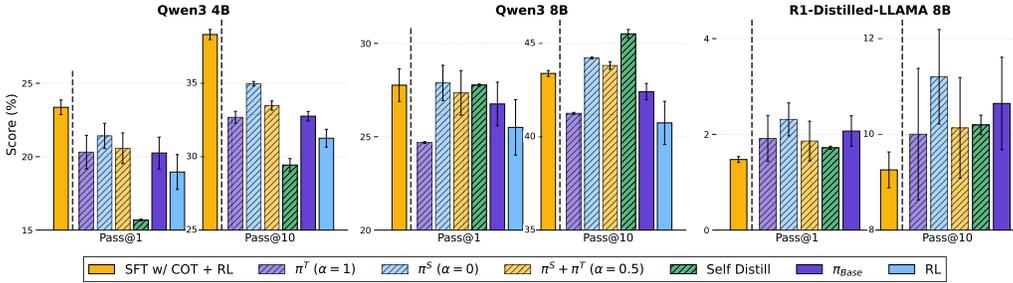


Figure 5: **Evaluation on Out-of-Domain Environments.** We report Pass@1 and Pass@10 on the GEM search-tool benchmark suite (7 datasets) for Qwen 3 models and R1-Distill-Llama-8B, using the best checkpoint selected on τ -Bench Retail. Bars show mean \pm standard errors over three seeds per dataset, comparing π -Distill variants (π^T ($\alpha = 1$), π^S ($\alpha = 0$), $\pi^T + \pi^S$ ($\alpha = 0.5$)) and OPSD against SFT w/ CoT + RL, π_{Base} , and standard RL.

E DETAILED ALGORITHMS

Here we outline the detailed algorithms for both π -Distill and OPSD.

Algorithm 1 π -Distill: Privileged Information Distillation for Language Models

- 1: **Input:** Dataset $\mathcal{D} = \{(s, \mathbf{I})\}$ where \mathbf{I} is privileged information, Initial Policy π_θ , Reference π_{ref} , α, β, ϵ , Learning rate η
 - 2: **Initialize:** $\phi \leftarrow \theta$ {Parameters are shared between Teacher and Student}
 - 3: **while** not converged **do**
 - 4: Sample batch $B = \{(s_i, \mathbf{I}_i)\}_{i=1}^N \sim \mathcal{D}$
 - 5: // **Step 1: Teacher Rollout (with Privileged Information)**
 - 6: **for each** $(s, \mathbf{I}) \in B$ **do**
 - 7: Sample K trajectories $\{\mathbf{o}_1, \dots, \mathbf{o}_K\} \sim \pi^T(\cdot | s, \mathbf{I})$
 - 8: Compute rewards $R(\mathbf{o}_k, s) = R_{env}(\mathbf{o}_k, s) - \beta D_{KL}[\pi^T(\cdot | s, \mathbf{I}) || \pi_{ref}(\cdot | s, \mathbf{I})]$
 - 9: **end for**
 - 10: // **Step 2: Compute Group-Centered Advantages**
 - 11: **for each** $k \in \{1, \dots, K\}$ **do**
 - 12: $\bar{R} = \frac{1}{K} \sum_{j=1}^K R(\mathbf{o}_j, s)$
 - 13: $A_k = R(\mathbf{o}_k, s) - \bar{R}$
 - 14: **end for**
 - 15: // **Step 3: Compute Objectives**
 - 16: **Teacher Objective (GRPO):**
 - 17: $J_{Teacher} = \frac{1}{K} \sum_{k=1}^K \min(\rho_k^{teacher} A_k, \text{clip}(\rho_k^{teacher}, 1 - \epsilon, 1 + \epsilon) A_k)$
 - 18: **Student Objective (Off-Policy GRPO):**
 - 19: Compute IS weights: $\rho_k = \frac{\pi^S(\mathbf{o}_k | s)}{\pi^T(\mathbf{o}_k | s, \mathbf{I})}$ {Note: Student input is s only}
 - 20: $J_{Student} = \frac{1}{K} \sum_{k=1}^K \min(\rho_k A_k, \text{clip}(\rho_k, 1 - \epsilon, 1 + \epsilon) A_k)$
 - 21: // **Step 4: Joint Update**
 - 22: $J_{\pi-Distill}(\theta) = \alpha J_{Teacher} + (1 - \alpha) J_{Student}$
 - 23: $\theta \leftarrow \theta + \eta \nabla_\theta J_{\pi-Distill}(\theta)$
 - 24: **end while**
-

Algorithm 2 On-Policy Self-Distillation

```

1: Input: Dataset  $\mathcal{D} = \{(s, \mathbf{I})\}$  where  $\mathbf{I}$  is privileged information, Initial Policy  $\pi_\theta, \beta, \epsilon$ ,
   Learning rate  $\eta$ 
2: Initialize:  $\theta$  {Parameters for the Student Policy}
3: while not converged do
4:   Sample batch  $B = \{(s_i, \mathbf{I}_i)\}_{i=1}^N \sim \mathcal{D}$ 
5:   // Step 1: Student Rollout (On-Policy sampling)
6:   for each  $(s, \mathbf{I}) \in B$  do
7:     Sample  $K$  trajectories  $\{\mathbf{o}_1, \dots, \mathbf{o}_K\} \sim \pi^S(\cdot | s)$ 
8:     {Reward Computation with Reverse KL}
9:     Compute rewards  $R(\mathbf{o}_k, s) = R_{\text{env}}(\mathbf{o}_k, s) - \beta D_{\text{KL}}[\pi^S(\cdot | s) || \pi^T(\cdot | s, \mathbf{I})]$ 
10:    end for
11:  // Step 2: Compute Group-Centered Advantages
12:  for each sample  $i \in \{1, \dots, N\}$  do
13:     $\bar{R}_i = \frac{1}{K} \sum_{j=1}^K R(\mathbf{o}_{i,j}, s_i)$ 
14:    for each  $k \in \{1, \dots, K\}$  do
15:       $A_{i,k} = R(\mathbf{o}_{i,k}, s_i) - \bar{R}_i$ 
16:    end for
17:  end for
18:  // Step 3: Objective (Off-Policy GRPO)
19:  Compute IS weights:  $\rho_{i,k} = \frac{\pi_\theta^S(\mathbf{o}_{i,k} | s_i)}{\pi_{\text{old}}^S(\mathbf{o}_{i,k} | s_i)}$ 
20:   $J(\theta) = \frac{1}{K} \sum_{k=1}^K \min(\rho_{i,k} A_{i,k}, \text{clip}(\rho_{i,k}, 1 - \epsilon, 1 + \epsilon) A_{i,k})$ 
21:  // Step 4: Policy Update
22:   $\theta \leftarrow \theta + \eta \nabla_\theta J(\theta)$ 
23: end while

```

F ADDITIONAL ANALYSIS

G OUT OF DOMAIN EXPERIMENTS (OOD)

In this section, we demonstrate that π -Distill generalizes effectively to OOD tasks, consistently outperforming standard RL and the base model, with this holding true for OPSD on **Qwen3-8B**. We report Pass@1 and Pass@10 metrics on the GEM (Liu et al., 2025b) search-tool benchmark suite which consists of 7 datasets. To simulate a realistic deployment scenario, we select the single best-performing checkpoint in τ -Bench retail for each model. We then evaluate these checkpoints across the suite using three random seeds and report the aggregated mean and standard errors in fig. 5.

We find that for both **Qwen3** models, π -Distill variants consistently outperform the base model and standard RL holding true for OPSD under more capable models. When compared to SFT w/ CoT + RL, we observe differences based on model size. On **Qwen3-4B**, SFT w/ CoT + RL is consistently the top performer with OPSD showing significant degradation. For **Qwen3-8B**, however, we find that $\alpha = 0$ and $\alpha = 0.5$ variants of π -Distill as well as OPSD can significantly outperform SFT w/ CoT + RL. For π -Distill, this aligns with our findings in section 6. Where for **Qwen3-4B** not all π -Distill variants show substantial improvements over SFT w/ CoT + RL, whereas for **Qwen3-8B**, the opposite trend is observed. We observe a similar trend for OPSD, where it shows significant degradation on **Qwen3-4B**, while on **Qwen3-8B** it shows significant improvements. We attribute this to the possibility that on smaller models OPSD may overfit the teacher supervision for the task, while the more potent reasoners can provide more generalizable feedback. Overall, results on both π -Distill and OPSD imply that explicit supervision from frontier model’s CoT is more important for smaller models, where stronger reasoners benefit from being more on-policy via self-generated CoT. Additionally, we observe deterioration in **R1-Distill-Llama-8B** under SFT w/ CoT + RL, where performance drops below the base model. While π -Distill and OPSD do not improve over the base model in this specific case, they avoid significant deterioration. Finally, raw RL consistently exhibits degradation relative to the base model

across all evaluations. Overall, these results indicate that π -Distill is highly effective for distillation in OOD settings when CoT traces are occluded, preventing performance from regressing below the base model. Moreover, both π -Distill and OPSD show significant improvements as model size scales.

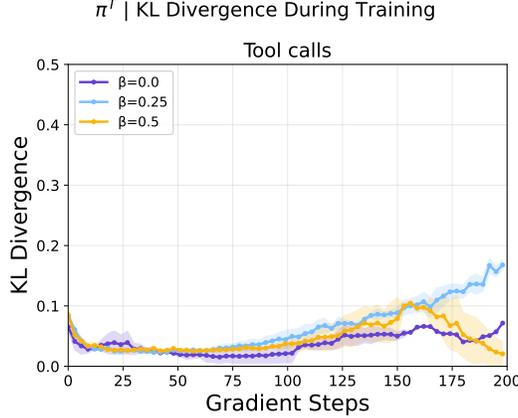


Figure 6: Training KL between π^T and π^S during training the teacher $\pi^S(\alpha = 1)$ on τ -Bench. We observe an early KL collapse making $\pi^S \approx \pi^T$. We attribute the underperformance of π^T on low KL settings to this collapse.

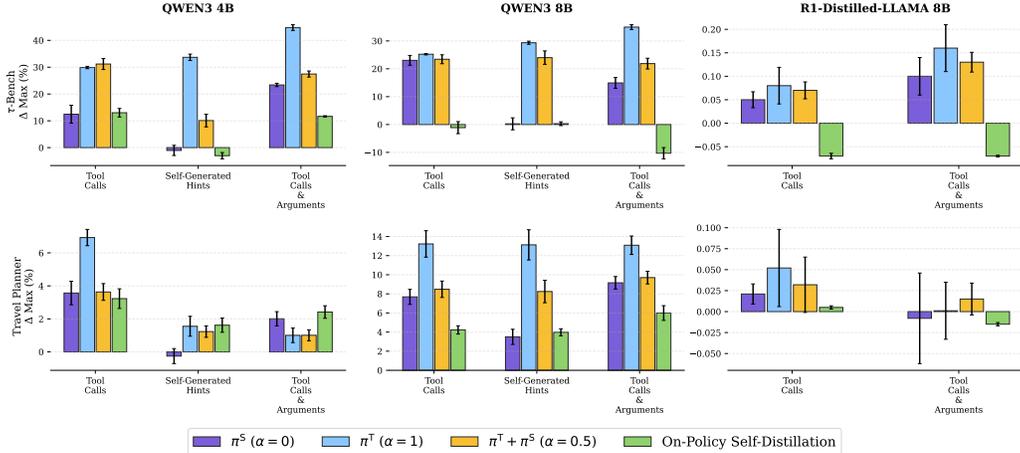


Figure 7: Maximum Improvement (Δ_{\max}) across PI Types and Algorithms. We compare the peak performance gain over baselines for different PI variants (x-axis) and training configurations (α) on TravelPlanner (bottom) and τ -Bench (top). We find that PI types which initially underperform (e.g., self-generated hints on Planner, Qwen3-8B) can yield substantial gains when the teacher is trained to utilize them ($\alpha > 0$), confirming that learning to leverage PI is an important factor in transferring from teacher π^T to student π^S .

G.1 WHAT MATTERS FOR OPSD

Figure 8 similarly to fig. 2 displays performance on τ -Bench retail and Travel Planner across PI types outlined in Section 5.2.

Findings. Our main finding is that, contrary to π -Distill, high-KL PI types do not necessarily indicate worse performance. Instead, the information content of the PI is the strongest predictor of successful transfer. For example, Tool Calls & Arguments, being the richest in

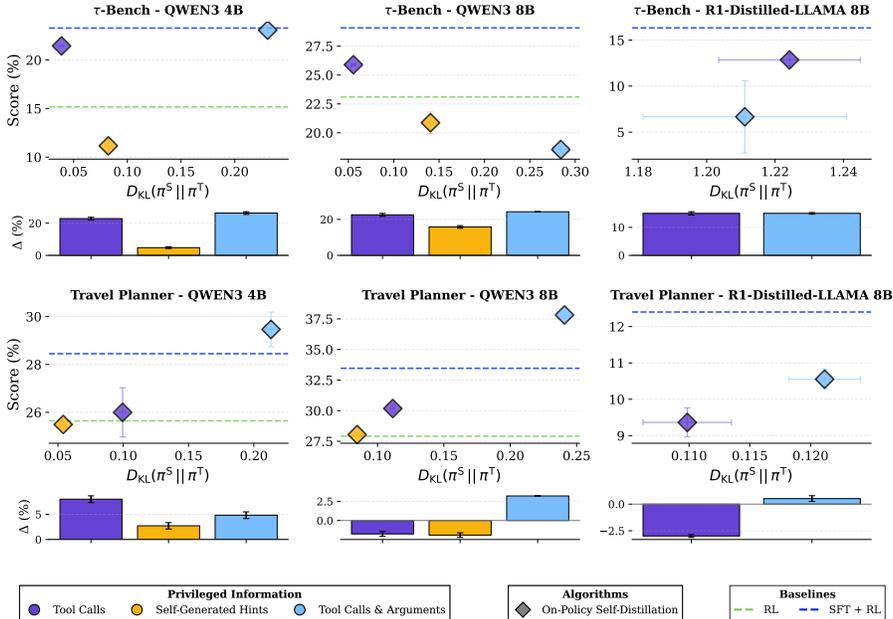


Figure 8: **Performance and Stability Analysis of OPSD.** We compare held-out performance on τ -Bench (top row) and Travel Planner (bottom row) across three base models and three PI types (colors). The scatter plots map final scores against the student-teacher KL divergence ($D_{KL}(\pi^S || \pi^T)$), while the bar-charts display the PI utility (Δ). Key observations: (1) Unlike π -Distill, higher KL are not always detrimental, rather information richness of PI is most important, finding (**Tool Calls & Arguments**) often performs best (e.g., all results on Travel Planner and Qwen3-4B on τ -Bench). (2) Excessive KL can override positive utility (note Qwen3-8B on τ -Bench, where $\Delta > 0$ but the high KL degrades performance). (3) R1-Distill-Llama-8B consistently struggles, which we attribute to either extreme KL divergence (τ -Bench) or negative PI utility (TravelPlanner).

information, consistently performs best for both Qwen3 models on Travel Planner, and also achieves the best performance for Qwen3-4B on τ -Bench.

Interestingly, although we find that high-information PI is generally best suited for transfer, we observe exceptions. In particular, for Qwen3-8B on τ -Bench, Tool Calls & Arguments performs the worst and exhibits the highest KL. In this case, we observe that Δ_{max} is negative, indicating that the reverse-KL penalty can inhibit training. Additionally, regarding the failure on R1-Distill-Llama-8B, while this could be due to the SFT phase prior to OPSD, we find that on τ -Bench the KL for both PI types is exceptionally high. While on Travel Planner, the PI utility, Δ is low or negative, possibly explaining the diminished results. Overall, we find that the strongest predictor of performance for OPSD is the information content of the PI. While using the richest form of PI can lead to the best results, careful consideration is required to ensure the KL does not become excessive and hinder training.

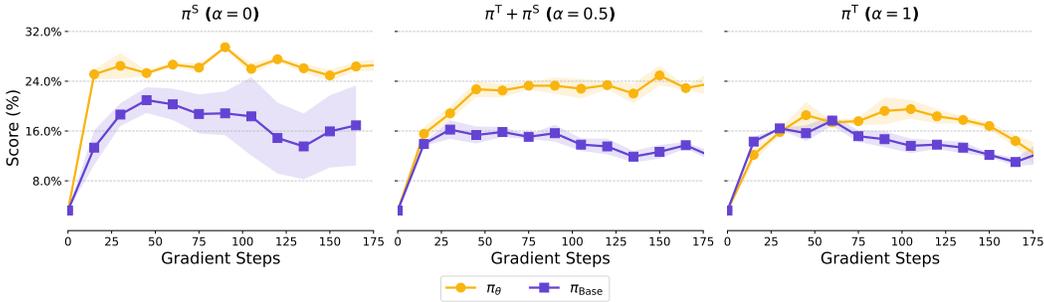


Figure 9: Ablation between using reverse-Kl between the teacher and a prior. We ablate over two possible priors, π_{base} and π_{θ}^S . We find that using π_{θ}^S to be highly important in obtaining best performance.

H ADDITIONAL ABLATIONS

H.1 PRIVILEGE INFORMATION LEAKAGE ABLATION

We measure privileged-information leakage with a simple keyword detector applied to each assistant message, using the keyword list [“privileged information”, “privileged info”, “priv info”, “secret information”, “secret info”, “correct tool calls”, “secret”, “privileged”, “hint”, “hints”]. For each occurrence we assign a penalty of -0.1 and accumulate it over the trajectory, and we mark a trajectory as leaking if any of these appear. Fig 10 reports results at $\beta = 0.25$ and shows that enabling this leakage penalty does not change task performance in any noticeable way, since the learning curves with and without the penalty overlap closely. Fig 11 shows the corresponding leakage rate during training, and it rises as training progresses across all modes. The pattern with tool-calls showing the highest leakage, tool calls and arguments increasing more moderately, and self-generated hints starting lower but still increasing over time.

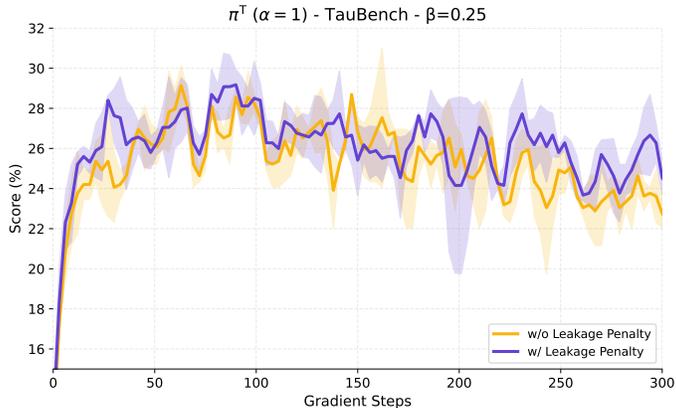


Figure 10: Performance with and without leakage penalty for $\pi^T(\alpha = 1)$. We find that although the penalty does reduce the leakage of the privileged information (see Figure 11) it does not affect performance.

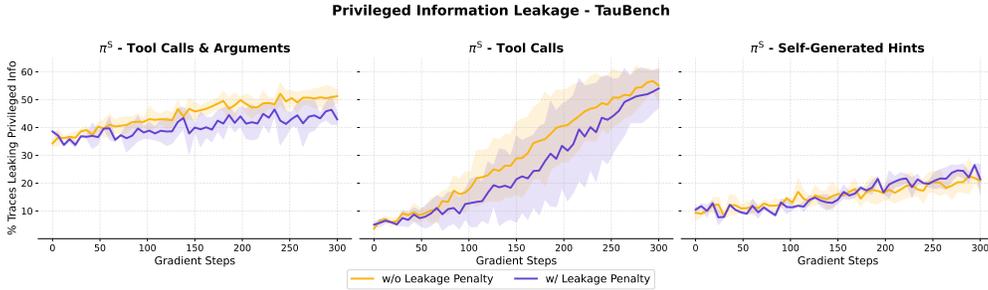


Figure 11: Proportion of traces leaking privileged information as training progresses, we see that regardless PI type, the leakage increases with more gradient steps, finding using a leakage penalty reduces this proportion, but not substantially.

Test-time leakage. As an additional analysis, we analyze whether leakage increases as models are trained with privileged data when evaluating π^S . Figure 12 outlines this for Qwen3-8B for 300 gradient steps. We empirically validate that leakage does not significantly affect or increase when evaluating using π^S as the presence of leaked words does not meaningfully increase as training continues.

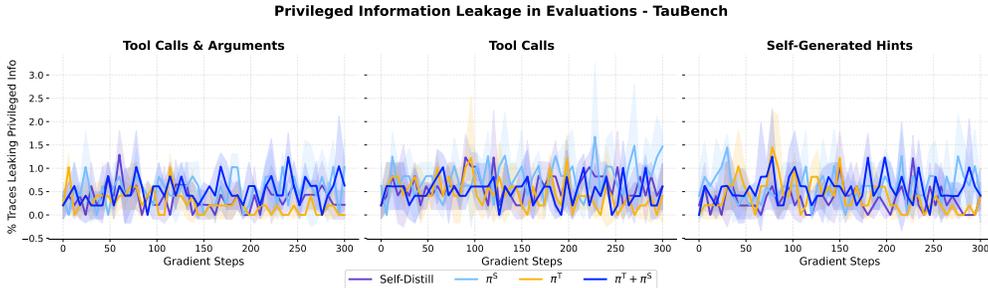


Figure 12: Proportion of traces leaking privileged information when **evaluating** π^S as training progresses, we see that there is no increase in leakage when evaluated as training goes on.

H.2 REFERENCE MODEL ABLATION

We ablate the choice of reference policy used in the KL term by comparing a fixed base reference, π_{base} , against using the student itself as the reference, π_θ (with stop-gradient on the reference branch). Fig.9 shows a clear difference in training stability. When π_{base} is used as the reference, the performance degrades over training and can collapse most clearly in the π^h setting ($\alpha = 1$), where the policy is pushed to move far from the base while still being penalized for that same deviation. In contrast, using π_θ as the reference (orange) yields stable learning across $\alpha \in \{0, 0.5, 1\}$, since the KL regularizer stays aligned with the current student distribution. Practically, this choice is also cheaper because using π_θ as the reference avoids maintaining a separate frozen reference model on the GPU, reducing memory and compute overhead.

I ABLATION ON β

In this section, we analyze β , the term controlling the regularization between the student and teacher, finding it important for achieving the best performance in 17/21 ablated configurations. We conduct extensive ablations on β keeping all other parameters fixed, sweeping over $\beta = \{0, 0.1, 0.25, 0.5\}$. We analyze *all* values of α and PI types for both π -Distill and OPSD on τ -Bench using both Qwen3 models, as well as the best-performing PI type for Travel Planner. All results are reported over three random seeds.

Figure 13 shows a subset of the learning curves on held-out data for the best-performing PI type and α value in τ -Bench Retail and Travel Planner, while Figure 14 provides the full set for τ -Bench. For π -Distill, we consistently find that $\beta > 0$ aids in obtaining the best performance, particularly when the teacher is being trained ($\alpha > 0$). While β can be sensitive, with no single value being the best in all settings, setting β higher than 0 generally allows for better or matching performance, with only four cases showing inferior results. For OPSD, we find that the value of β is less important with our results on Section 7 showing that information granularity and student-teacher KL are more important factors.

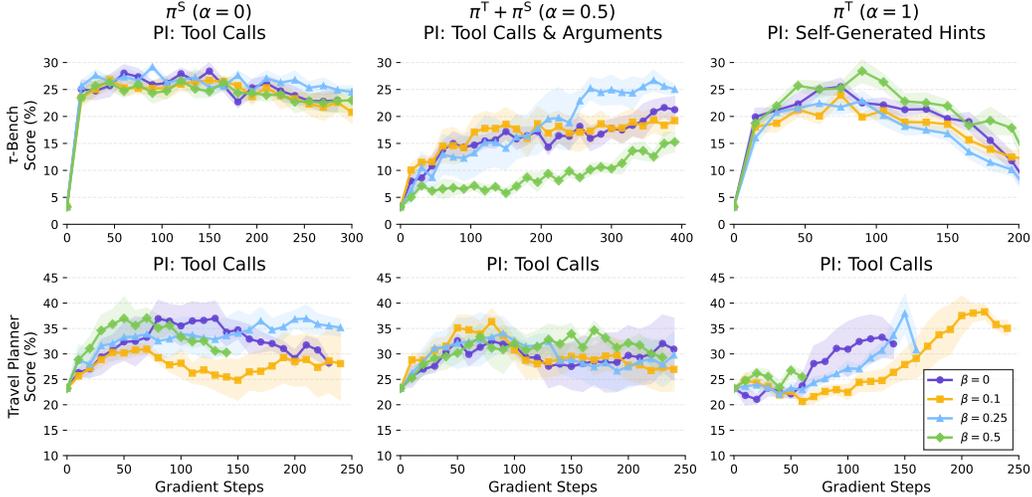


Figure 13: **Evaluation performance throughout training for π -Distill variants across varying KL penalties (β).** Runs that deteriorate significantly early are truncated for visual clarity. We observe that for settings involving teacher updates ($\alpha > 0$), a non-zero penalty ($\beta > 0$) is crucial for stabilizing training and achieving peak performance. Here error bars indicate standard errors. Discrepancies between plot and table values are addressed in App. K

I.1 FULL β ABLATIONS

Figure 14 outlines the full suite of β ablations for π -Distill variants and self-distillation for Qwen3-4/8B. We sweep over $\beta \in \{0.1, 0.25, 0.5\}$.

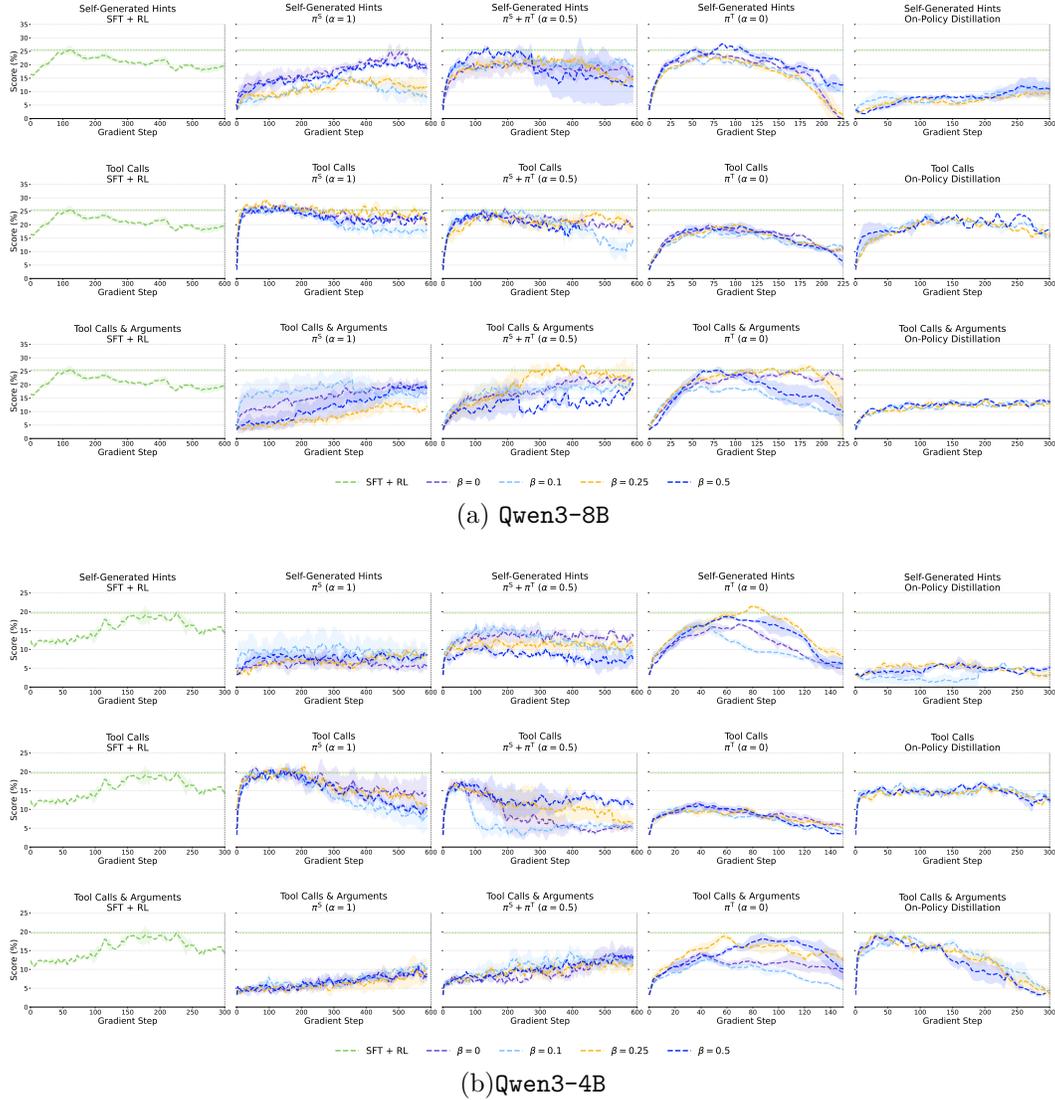


Figure 14: Full sweep for β over variants of π -Distill. All experiments are given 600 gradient step budget, where here we cutoff experiment types early if they crash or do not continue learning. Consistent performance gains are seen for $\beta > 0$ across both 8B and 4B scales. We find that $\beta > 0$ is important in 14/18 cases. Specifically, we find it most important when training π^T and less important with student-only training π^S .

J IMPLEMENTATION SETUP - FURTHER DETAILS

Implementation Details. We run all experiments using 2 H100 GPUs with a context limit of 25k tokens (with one exception, see App. J). We found initially that traces often exceeded this limit, thus adopting a length penalty reward that penalizes trajectories exceeding 15k tokens (full details in App. J.1). Additionally, we observed during training that a few tokens referencing the PI consistently exhibited very high KL (e.g., the token “hint”). We thus incorporate a penalty on the frequency of these tokens. Our final experiments incorporate this penalty, though we found in practice it makes little difference for final performance (see App. H for full details). All PI is added to the task description in the system prompt, using the prompts in App. L.4. We thoroughly sweep over relevant HPs for all baselines, for a full list see App. J.2. Using the best performing set, we run three seeds of this set and report their average. In experiments, we score models using the best running average over three subsequent checkpoints, capturing both performance and stability.

J.1 LENGTH PENALTY

In our setting, we found that policies can become overly verbose. To mitigate this, we add a cosine-shaped length penalty inspired by the cosine length-scaling reward Yeo et al. (2025). We apply this penalty only to successful traces (i.e., when the base reward $r > 0$). For each assistant turn i with token length l_i , we use a no-penalty threshold $l_{\text{th}} = 2000$ and a soft allowance $l_{\text{max}} = 5000$. If $l_i \leq l_{\text{th}}$, the turn penalty is $p_i = 0$ otherwise we assign a negative penalty using a linear-cosine schedule that increases in magnitude with length, approaching $-\lambda$ near l_{max} with $\lambda = 0.1$, and becoming harsher beyond l_{max} (up to an endpoint of -2λ). We then average the per-turn penalties across the N assistant turns:

$$\bar{p} = \frac{1}{N} \sum_{i=1}^N p_i.$$

We cap the total penalty as $p_{\text{total}} = \max(\bar{p}, -0.3)$ and add that to the base reward $r' = r + p_{\text{total}}$.

J.2 HYPERPARAMETER CHOICES

We run an extensive sweep over hyper-parameters for all baselines and our methods. We summarize hyperparameters used for each model in Table 2.

Category	Parameter	R1-Distill-Llama-8B	QWEN3-4B	QWEN3-8B
General	Seeds	3	3	3
	Rollout temperature	0.75	0.75	0.75
	Trace length filter	Discard if tokens > 25k (TauBench/TravelPlanner), > 35k (RL/OPSD)		
	Advantage processing	pop zero-advantage (always)		
Training Budgets	TauBench total gradient steps	600	600	600
	TravelPlanner total gradient steps	400	400	400
Sampling	Gradient steps per sampling	TauBench = 3, TravelPlanner = 2		
	Repeats per group	TauBench = 5, TravelPlanner = 4		
	Training tasks sampled	TauBench = 64, TravelPlanner = 45, SFT+RL (TauBench) = 128		
Optimization	Learning-rate sweep	{1e-6, 5e-6, 1e-5}		
	Final LR	TauBench = 5e-6, π -distill = 1e-5, RL/OPSD = 5e-6		
	β	TravelPlanner/OPSD = 0.5, π -distill = 0.25 (unless swept)		
	Clipping/epsilon	Lower = 0.8, Upper = 1.2		
Alpha Annealing	$\alpha = 0.5$ schedule	Linearly anneal $\alpha : 0 \rightarrow 0.5$ over 15 epochs		

Table 2: Hyperparameters used for each fine-tuned model. All models share identical hyperparameter settings.

K PLOT AND TABLE DISCREPANCIES

In this Fig 15, the 31.11% number is computed from the three seed peaks in the right panel. Each seed reaches its best score at a different gradient step, roughly one around 70 steps, another around 95, and another around 150. We take those three peak values and average them, which gives 31.11%, and that is what is reported in the table. The left panel instead shows the mean score at each gradient step, averaging the seeds at the same step, so its peak corresponds to the best point of the averaged curve rather than the average of the three best points. This difference in aggregation means the table can be higher than the peak of the averaged curve shown in the left plot.

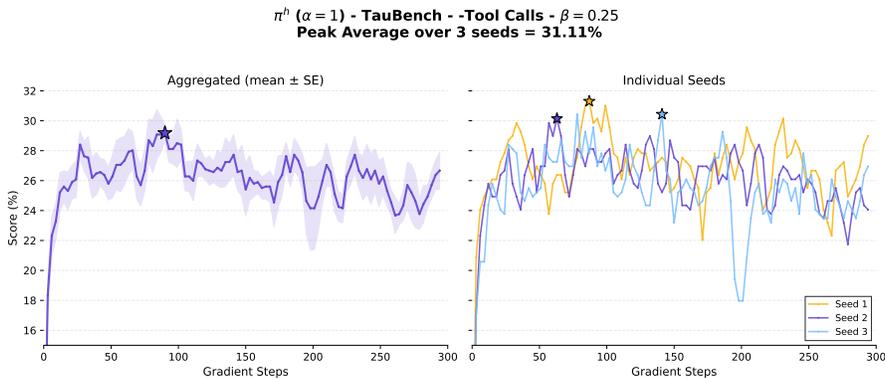


Figure 15: Comparison aggregating training runs versus the individual training runs. The left figure displays the mean score across seeds at each gradient step, while the right panel highlights the individual trajectories where peaks occur at different intervals. The reported table value (31.11%) represents the average of these individual seed peaks.

L PROMPTS

L.1 SYSTEM PROMPTS

For both τ -Bench and Travel Planner, we slightly modify the prompt to indicate that the model should output its action in corresponding `<action></action>` tags. Below, we display the full system prompts used for τ -Bench and Travel Planner.

τ -Bench Retail System Prompt

Retail agent policy

As a retail agent, you can help users cancel or modify pending orders, return or exchange delivered orders, modify their default user address, or provide information about their own profile, orders, and related products.

- At the beginning of the conversation, you have to authenticate the user identity by locating their user id via email, or via name + zip code. This has to be done even when the user already provides the user id.
- Once the user has been authenticated, you can provide the user with information about order, product, profile information, e.g. help the user look up order id.
- You can only help one user per conversation (but you can handle multiple requests from the same user), and must deny any requests for tasks related to any other user.
- Before taking consequential actions that update the database (cancel, modify, return, exchange), you have to list the action detail and obtain explicit user confirmation (yes) to proceed.
- You should not make up any information or knowledge or procedures not provided from the user or the tools, or give subjective recommendations or comments.
- You should at most make one tool call at a time, and if you take a tool call, you should not respond to the user at the same time. If you respond to the user, you should not make a tool call.

Domain basic

- All times in the database are EST and 24 hour based. For example “02:30:00” means 2:30 AM EST.
- Each user has a profile of its email, default address, user id, and payment methods. Each payment method is either a gift card, a paypal account, or a credit card.
- Our retail store has 50 types of products. For each type of product, there are variant items of different options. For example, for a ‘t shirt’ product, there could be an item with option ‘color blue size M’, and another item with option ‘color red size L’.
- Each product has an unique product id, and each item has an unique item id. They have no relations and should not be confused.
- Each order can be in status ‘pending’, ‘processed’, ‘delivered’, or ‘cancelled’. Generally, you can only take action on pending or delivered orders.
- Exchange or modify order tools can only be called once. Be sure that all items to be changed are collected into a list before making the tool call!!!

Cancel pending order

- An order can only be cancelled if its status is ‘pending’, and you should check its status before taking the action.
- The user needs to confirm the order id and the reason (either ‘no longer needed’ or ‘ordered by mistake’) for cancellation.

- After user confirmation, the order status will be changed to ‘cancelled’, and the total will be refunded via the original payment method immediately if it is gift card, otherwise in 5 to 7 business days.

Modify pending order

- An order can only be modified if its status is ‘pending’, and you should check its status before taking the action.
- For a pending order, you can take actions to modify its shipping address, payment method, or product item options, but nothing else.

Modify payment

- The user can only choose a single payment method different from the original payment method.
- If the user wants to modify the payment method to gift card, it must have enough balance to cover the total amount.
- After user confirmation, the order status will be kept ‘pending’. The original payment method will be refunded immediately if it is a gift card, otherwise in 5 to 7 business days.

Modify items

- This action can only be called once, and will change the order status to ‘pending (items modified)’, and the agent will not be able to modify or cancel the order anymore. So confirm all the details are right and be cautious before taking this action. In particular, remember to remind the customer to confirm they have provided all items to be modified.
- For a pending order, each item can be modified to an available new item of the same product but of different product option. There cannot be any change of product types, e.g. modify shirt to shoe.
- The user must provide a payment method to pay or receive refund of the price difference. If the user provides a gift card, it must have enough balance to cover the price difference.

Return delivered order

- An order can only be returned if its status is ‘delivered’, and you should check its status before taking the action.
- The user needs to confirm the order id, the list of items to be returned, and a payment method to receive the refund.
- The refund must either go to the original payment method, or an existing gift card.
- After user confirmation, the order status will be changed to ‘return requested’, and the user will receive an email regarding how to return items.

Exchange delivered order

- An order can only be exchanged if its status is ‘delivered’, and you should check its status before taking the action. In particular, remember to remind the customer to confirm they have provided all items to be exchanged.
- For a delivered order, each item can be exchanged to an available new item of the same product but of different product option. There cannot be any change of product types, e.g. modify shirt to shoe.
- The user must provide a payment method to pay or receive refund of the price difference. If the user provides a gift card, it must have enough balance to cover the price difference.
- After user confirmation, the order status will be changed to ‘exchange requested’, and the user will receive an email regarding how to return items. There is no need to place a new order.

Available tools:

```
{tools_info}
```

Instruction

You need to act as an agent that use the above tools to help the user according to the above policy. At each step, your generation should have exactly the following format:

```
<think>
...Few lines of reasoning
</think>

<action>
{"name": <The name of the action>, "arguments": <The arguments to the action in
  json format>}
</action>
```

The Action will be parsed, so it must be valid JSON and within the `<action>` and `</action>` tags. You should not use made-up or placeholder arguments.

For example, if the user says “I want to know the current weather of San Francisco”, and there is such a tool available:

```
{
  "type": "function",
  "function": {
    "name": "get_current_weather",
    "description": "Get the current weather",
    "parameters": {
      "type": "object",
      "properties": {
        "location": {
          "type": "string",
          "description": "The city and state, e.g. San Francisco, CA",
        },
        "format": {
          "type": "string",
          "enum": ["celsius", "fahrenheit"],
          "description": "The temperature unit to use. Infer from
            location.",
        },
      },
      "required": ["location", "format"],
    },
  },
}
```

Example response**Step 1:**

```
<think>
... Few lines of reasoning
</think>

<action>
{"name": "get_current_weather", "arguments": {"location": "San Francisco, CA",
  "format": "fahrenheit"}}
</action>
```

The tool and the user have the same id tags so if the user returns “70F”, your response can be:

Step 2:

```
<think>
... Few lines of reasoning
```

```
</think>
```

```
<action>
```

```
{"name": RESPOND_ACTION_NAME, "arguments": {"RESPOND_ACTION_FIELD_NAME": "The
  current weather of San Francisco is 70F."}}
```

```
</action>
```

Requirement

Try to be helpful and always follow the policy. Always try to validate your steps in your thinking and checkover your work, try to predict what will happen given your actions. Always make sure you generate valid JSON only.

Important Notes

- Be very brief in the reasoning, do not repeat the entire context or the tools, just focus on what you need to do next.
- Always respond to the user using the tool {RESPOND_ACTION_NAME} to ensure the user sees your response.
- Only use a single set of thinking tags, the user cannot see your thoughts.
- Always wrap your tool calls <action> and </action> tags or else the system will not be able to parse your actions.
- Always start your outputs by thinking using <think> ... </think>.

Ok with this said, let us reason this out step by step always starting with <think> and ending with </think>.

Travel Planning System Prompt

You are a travel planning assistant. You have access to the following tools:

- FlightSearch[Departure City, Destination City, Date]
- AccommodationSearch[City]
- RestaurantSearch[City]
- AttractionSearch[City]
- CitySearch[State]
- GoogleDistanceMatrix[Origin, Destination, Mode]
- NotebookWrite[Description]
- Planner[Query]

Respond in this format:

```
<think>Your reasoning here</think>
```

```
<action>ToolName[arguments]</action>
```

Query:

```
{query}
```

L.2 FULL TRAJECTORIES

Example Trajectory

User Task:

I'd like to start a return for #W6619432 and process it through PayPal account paypal_3738584 for the Dumbbell Set and Yoga Mat.

Agent Trajectory:

1. Turn 1:

```

<action>
{"name": "respond", "arguments": {"content": "I'd be happy to help you
  with your return for order #W6619432. First, I need to verify your
  identity to ensure the security of your account. Could you please
  provide your email address? Alternatively, you can provide your
  first name, last name, and zip code."}}
</action>

2. Turn 2:

<action>
{"name": "find_user_id_by_name_zip", "arguments": {"first_name":
  "Sophia", "last_name": "Nguyen", "zip": "20171"}}
</action>

3. Turn 3:

<action>
{"name": "get_order_details", "arguments": {"order_id": "#W6619432"}}
</action>

4. Turn 4:

<action>
{"name": "respond", "arguments": {"content": "Thank you for providing
  your information. I've verified your order #W6619432, which is
  delivered and includes:\n- Yoga Mat (item ID: 6195938807)\n-
  Dumbbell Set (item ID: 3735133539)\n\nYou're requesting to return
  both items and process the refund through your PayPal account
  (paypal_3738584). The total refund amount will be
  $612.35.\n\nPlease confirm that you want to proceed with returning
  both items and receiving the refund to your PayPal account. Type
  'yes' to confirm or 'no' to cancel."}}
</action>

5. Turn 5:

<action>
{"name": "return_delivered_order_items", "arguments": {"order_id":
  "#W6619432", "item_ids": ["6195938807", "3735133539"],
  "payment_method_id": "paypal_3738584"}}
</action>

6. Turn 6:

<action>
{"name": "respond", "arguments": {"content": "Your return request for
  order #W6619432 has been processed successfully. The order status
  is now 'return requested'. You will receive an email with detailed
  instructions on how and where to return the Yoga Mat and Dumbbell
  Set. The refund of $612.35 will be processed to your PayPal account
  (paypal_3738584). Please check your email for the return
  instructions. Is there anything else I can help you with today?"}}
</action>

7. Turn 7:

<action>
{"name": "respond", "arguments": {"content": "You're very welcome! I'm
  glad I could help with your return. Your refund will be processed
  to your PayPal account. Have a great day!"}}
</action>

```

L.3 SELF-GENERATED HINTS

Self-Generated Hint Example

Task:

Your name is Amelia Wilson and your email is `amelia.wilson1598@example.com`. You are confident, cautious, dependent, shy, pessimistic. Cancel order `#W3062096` because ordered by mistake.

Self-Generated Hint:

The agent authenticated the user via email, retrieved the order details, confirmed the order ID and cancellation reason, and successfully cancelled the order using the `cancel_pending_order` function. The refund process was initiated, and the user was informed of the next steps.

Conversation:

[user] Hi, I need to cancel an order I placed by mistake.

[assistant] Action: `{"name": "respond", "arguments": {"content": "I can help you cancel your order. First, I need to authenticate your identity. Could you please provide your email address? Alternatively, you can provide your first name, last name, and zip code."}}`

[user] Sure, my email is `amelia.wilson1598@example.com`.

[assistant] Action: `{"name": "find_user_id_by_email", "arguments": {"email": "amelia.wilson1598@example.com"}}`

[API output] `amelia_wilson_4614`

[assistant] Action: `{"name": "respond", "arguments": {"content": "Thank you for providing your email. Could you please share the order ID you'd like to cancel? This will help me verify the order details and proceed with the cancellation."}}`

[user] Yes, the order ID is `W3062096`.

[assistant] Action: `{"name": "get_order_details", "arguments": {"order_id": "#W3062096"}}`

[API output] `{"order_id": "#W3062096", "user_id": "amelia_wilson_4614", "address": {"address1": "388 Elm Avenue", "address2": "Suite 384", "city": "Dallas", "country": "USA", "state": "TX", "zip": "75215"}, "items": [{"name": "T-Shirt", "product_id": "9523456873", "item_id": "3234800602", "price": 46.66, "options": {"color": "red", "size": "L", "material": "cotton", "style": "v-neck"}}, {"name": "Smart Watch", "product_id": "6945232052", "item_id": "8739626972", "price": 370.87, "options": {"color": "silver", "band material": "silicone", "display": "AMOLED"}}, {"name": "Jigsaw Puzzle", "product_id": "1808611083", "item_id": "9370300555", "price": 45.9, "options": {"pieces": "1000", "theme": "art", "difficulty level": "expert"}}, {"fulfillments": [], "status": "pending", "payment_history": [{"transaction_type": "payment", "amount": 463.43, "payment_method_id": "paypal_4101143"}]}`

[assistant] Action: `{"name": "respond", "arguments": {"content": "I can cancel your order #W3062096, which includes a T-Shirt, Smart Watch, and Jigsaw Puzzle, totaling $463.43. The reason for cancellation is 'ordered by mistake'. The refund will be processed`

```
back to your PayPal account in 5-7 business days. Please confirm with 'yes' if you'd like to proceed with the cancellation."}]}
```

```
[user] Yes, please go ahead and cancel it.
```

```
[assistant] Action: {"name": "cancel_pending_order",
"arguments": {"order_id": "#W3062096",
"reason": "ordered by mistake"}}
```

```
[API output] {"order_id": "#W3062096",
"user_id": "amelia_wilson_4614", "address": {"address1":
"388 Elm Avenue", "address2": "Suite 384", "city": "Dallas",
"country": "USA", "state": "TX", "zip": "75215"}, "items":
[{"name": "T-Shirt", "product_id": "9523456873",
"item_id": "3234800602", "price": 46.66, "options":
{"color": "red", "size": "L", "material": "cotton",
"style": "v-neck"}}, {"name": "Smart Watch",
"product_id": "6945232052", "item_id": "8739626972",
"price": 370.87, "options": {"color": "silver",
"band material": "silicone", "display": "AMOLED"}},
{"name": "Jigsaw Puzzle", "product_id": "1808611083",
"item_id": "9370300555", "price": 45.9, "options":
{"pieces": "1000", "theme": "art", "difficulty level":
"expert"}]}, "fulfillments": [], "status": "cancelled",
"payment_history": [{"transaction_type": "payment",
"amount": 463.43, "payment_method_id": "paypal_4101143"},
{"transaction_type": "refund", "amount": 463.43,
"payment_method_id": "paypal_4101143"}],
"cancel_reason": "ordered by mistake"}
```

```
[assistant] Action: {"name": "respond", "arguments": {"content":
"Your order #W3062096 has been successfully cancelled. The refund
of $463.43 will be processed back to your PayPal account within
5-7 business days. If you have any further questions or need
assistance with anything else, feel free to ask!"}}
```

```
[user] Thank you, that's all I needed. ###STOP###
```

Hint Generation Prompt:

You are creating a concise natural language hint that will help an agent solve the task. Write a helpful hint describing the key actions and strategy the agent used. If the agent did not fully complete the task, focus on what they did achieve. **IMPORTANT:** Output **ONLY** plain text - no JSON, no code blocks, no markdown formatting. Keep the hint to 2-3 sentences maximum.

The prompt used

L.4 PRIVILEGED INFORMATION PROMPTS

We use the same prompts When using PI in both models. Specifically, we try to ask the model not to reveal it has access to PI by encasing it in `<Secret information></Secret information>` tags. We found this to be unsuccessful as shown in the experiments in App H.

Privileged Instructions

`<Secret information>`

This section contains privileged hints derived from prior successful trajectories. Use them to stay on track, but never mention that hints were provided.

Hint Summary:

`{hint_lines}`

Your Task:

1. Follow the hints to accomplish the task still reasoning with `<think>...</think>` before actions.
2. Never acknowledge having access to secret hints.

Always start responses with a `<think>` block.
`</Secret information>`

Privileged Tool Instructions

`<Secret information>`

This section is highly important for you to pay attention to. You have been given access to secret information that is not available to the user and should not be mentioned. Here is a successful set of tools, given the context and tools that have been called so far use this information to accomplish the task.

`{tool_calls_section}`

Your Task:

1. **You will need to use the provided tools in order to accomplish the task.**
2. **Reasoning:** Please reason about your steps. If you think you need to use a tool, provide reasoning justifications and clear logic for using it before invoking it using `<think>reasoning...</think>`.
3. **Strict Constraint:** Do NOT mention that you have been given access to the secret information. You will be penalized for violating this rule.

Reasoning Format

Your reasoning traces will be used to train future agents that do not have access to privileged information. Please make sure your reasoning is clear and concise. Always start your outputs by thinking using `<think> ...</think>`
`</Secret information>`

M REWARD HACKING IN TRAVEL PLANNER

Our early experiments show that the model is able to consistently reward-hack the original rubric-based rewards proposed by Xie et al. (2024). Specifically, Figure 16 shows the learning curves under the original reward structure. The policy appears to steadily improve; however, upon inspection, we observe that all traces collapse to the response structure shown in Appendix M. We attribute this behavior to the model learning that invoking the `Planner` tool ends the conversation and, with certain arguments, reliably yields a high reward.

We address this issue by removing the requirement that all “easy” constraints must be satisfied before any “hard” constraint is checked. Instead, we evaluate each easy constraint individually and, only if it is satisfied, we then check its corresponding hard constraint.

τ -Bench Travel Planner Model Trajectory

Role: Assistant

`<think>`

Okay, let’s tackle this travel planning query. The user wants a 7-day plan from Punta Gorda to three Wisconsin cities between March 16-22, 2022, with a \$5,700 budget.

First, I need to use the `Planner` tool to structure the itinerary. Check `FlightSearch` for available flights from Punta Gorda to Wisconsin cities. Maybe Milwaukee, Madison, and Green Bay? Use `AccommodationSearch` for each city. `RestaurantSearch` and `AttractionSearch` to suggest activities. Check `GoogleDistanceMatrix` for travel times between cities. Make sure the total costs are under budget. Use `NotebookWrite` to document the plan. Ensure the

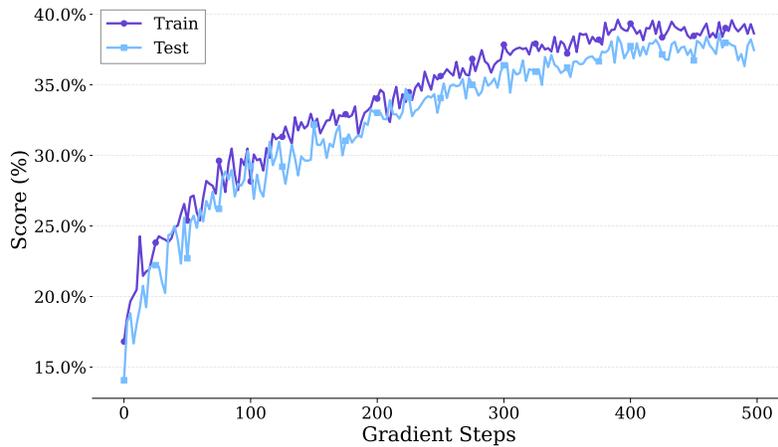


Figure 16: Training and testing curves for Travel Planner, using the original reward structure proposed in Xie et al. (2024) with Qwen3-4B. We find that while the model learns, the sequences generally converge on reward hacking. Where the entire score is dominated by the “easy constraint score” as the model has learnt to satisfy some basic constraints, leading to the policy outputting the same response for all queries.

```

Planner incorporates all elements and stays within the budget. Avoid
mentioning hints, just present the plan logically.
</think>

```

```

<action>Planner[Query: Craft a 7-day Wisconsin itinerary from Punta Gorda,
including 3 cities, flights, accommodations, attractions, and budget
adherence using available tools. Prioritize March 16-22, 2022.]</action>

```