
e3: Learning to Explore Enables Extrapolation of Test-Time Compute for LLMs

Amrith Setlur^{1,*} Matthew Y. R. Yang^{1,*} Charlie Snell² Jeremy Greer³
Ian Wu¹ Virginia Smith¹ Max Simchowitz¹ Aviral Kumar¹
¹Carnegie Mellon University ²UC Berkeley ³Oumi ^{*}Equal contribution

Abstract

Test-time scaling offers a promising path to improve LLM reasoning by utilizing more compute at inference time; however, the true promise of this paradigm lies in *extrapolation* (i.e., performance continues to improve on hard problems as LLMs keep “thinking” for longer, much beyond the maximum token budget they were trained on). Surprisingly, we find that most existing reasoning models do not extrapolate. We show that one way to enable extrapolation is by training the LLM at *in-context exploration*; that is, training the LLM to effectively spend its test time budget by chaining operations (such as generation, verification, refinement, etc.), and testing multiple hypotheses before it can commit to an answer. To enable in-context exploration, we identify three key ingredients as part of our recipe e3: (1) chaining asymmetries in base LLM competence, e.g., chaining verification (easy) with generation (hard), as a way to implement in-context search; (2) leverage negative gradients from incorrect traces to amplify exploration that chains additional asymmetries, resulting in longer search traces during RL; and (3) align task difficulty with training token budget to structure in-context exploration. Our recipe e3 produces the best known 1.7B model according to AIME/HMMT’25 scores, which can also extrapolate compute to $2.5\times$ the model training budget.

1 Introduction

Test-time scaling boosts large language model (LLM) performance by extending inference, spending more compute on “thinking” before producing an answer. Its ultimate promise lies in enabling models to continue improving performance by scaling test-time compute post-deployment. E.g., if the model can learn to implement “algorithmic procedures” like planning, self-reflection, or backtracking, it can discover more accurate responses as more test compute is used. With this motivation, current recipes post-train LLMs via reinforcement learning (RL) [1, 2] and supervised fine-tuning (SFT) [3, 4] at long context windows. However, it is unclear whether the models post-trained with current recipes can truly realize the promise of *extrapolation*: if we scale the test compute beyond the maximum *training budget*, would the LLM be able to continue to solve more and more problems?

Although performance at very long response lengths may be restricted by other factors like model architecture or model size [5], one can at least expect that an LLM should benefit from test-time scaling within the pretraining and fine-tuning training budget. Mechanistically, this can be realized by implementing algorithmic procedures (e.g., generate-verify-revise, best-of- N) in-context [6–9]. However, similar to other empirical studies of reasoning models [10, 11], we note that many open models perform poorly when extrapolating to $2\text{--}3\times$ the training budget (see App. B). Thus, relying on current RL/SFT recipes to yield effective extrapolation is mostly futile.

In this paper, we show that the key to enabling extrapolation is *learning to explore in-context*: if a model learns to use compute by searching through multiple reasoning paths or implementing algorithmic procedures, it can “guide” the search towards the correct answer, and improve its perfor-

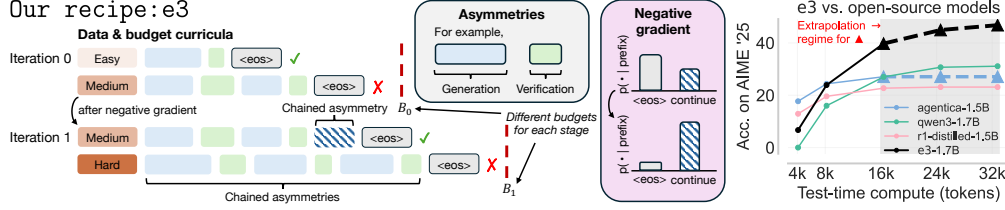


Figure 1: *In-context exploration enables extrapolation (e3)*: (i) chaining asymmetric capabilities in the base model, *e.g.*, reliably self-verifying responses after generating them; (ii) using negative gradients in RL training to penalize incorrectly terminated model responses, lengthening them further with more chained asymmetries, until the correct answer is discovered; and (iii) data & budget curricula for RL training that carefully balances explore-exploit tradeoff by sequentially training models on different datasets and training compute budgets. Qwen3-1.7B fine-tuned with e3 extrapolates test-time compute outperforming all $\leq 2\text{B}$ models on AIME’25.

mance as more test compute becomes available. Even under the original training compute budget, we expect learning to explore in-context to improve performance on unseen, out-of-distribution problems [12, 13]. To demonstrate this, we build a recipe e3, which trains models that leverage test compute for in-context exploration and can perform well at both normal training and extrapolation budgets. At its core, e3 is based on the following three ingredients (see Fig. 1):

1) Asymmetries are critical for learning to explore. LLMs can learn to explore only when each segment in the output trace is useful in “guiding” subsequent ones, *e.g.*, if verifying initial segments can lead to more refined answers later on. In the absence of external tools, we show that feedback can emerge from *asymmetries*, which are differences in the model’s competence at different procedures that constitute an output trace. In the context of self-verification, this corresponds to a verification-generation (VG) gap, where models are more capable of verifying their answers than generating correct ones. While prior work [14–18] has noticed such asymmetries, we show that these are critical for extrapolation, meaning that in their absence, scaling is strikingly hard. This is conceptualized by a didactic model of in-context exploration we study, termed p^k . Say an LLM is a perfect verifier with a failure probability of p on a best-guess response. If it learns to make k sequential attempts, verifying each time and stopping when it arrives at the correct answer, then the failure probability drops to p^k .

2) Negative gradient in RL amplifies in-context exploration. If asymmetries are a prerequisite for learning to explore, what enables them to evolve and facilitate learning useful exploration strategies during post-training? We show that *negative gradients* (*i.e.*, gradients on incorrect traces, see App. B) in RL training is a key enabler of in-context exploration, when the base model presents asymmetries. Negative gradients drive exploration [19, 20] by moving the probability mass from shorter failed traces onto longer traces that “chain” new asymmetries (*e.g.*, LLM verifying a calculation one more time). In contrast, SFT only maximizes likelihood on correct traces in the training data and reinforces the model to end the solution within the length of these traces. Conceptually, SFT only aims to reduce the failure probability p at a fixed k in our p^k model, whereas negative gradients also amplify k .

3) Structured exploration with coupled curriculum. Finally, while negative gradients amplify asymmetries and produce longer responses, running RL training at longer budgets suffers from poor training convergence, typically seen in long-horizon RL [21]. While one could train with a smaller budget, we show that training on hard problems at short context lengths often disincentivizes exploration altogether since the model is forced to commit to an answer prematurely. To resolve this, we design a *coupled curriculum* over pairs of (data mixture, training budget) that effectively structures the exploration driven by the negative gradient. Our key insight is that at any stage of the curriculum, we choose the smallest “RL optimization friendly” budget such that the model initialized for RL training can: (i) complete most of its responses within the budget; and (ii) can continue to improve performance as it chains more asymmetries beyond the chosen budget.

The above constitutes our recipe e3, which we use to post-train Qwen3-1.7B on a 16k budget. We achieve the **best performance at $< 2\text{B}$ scale on AIME’25 and HMMT’25**, and our model’s performance consistently improves as we extrapolate the test-time budget to 40k. We include a discussion of related work in App. A.

2 Asymmetries in the Base Model: A Prerequisite for In-Context Exploration

How can extrapolating beyond the training budget improve performance? To understand this, we begin by revisiting why longer traces perform better in general. The conventional wisdom is that

longer traces can represent solutions that make multiple attempts, interleaving verification and generation [14, 22, 23], to arrive at the final answer. We can think of this as the LLM learning to interleave basic “skills”, *e.g.*, verification, summarization, or retrieval, to perform in-context exploration. But why, or when, should a post-training recipe favor learning such solution traces over other uses of test-time compute that arrive at an answer more directly? This section demonstrates that when the base model exhibits *asymmetric* incompetence at different skills, RL post-training prefers to learn solutions that *chains asymmetric skills* in ways that improve final performance (see Definition 2.1 for a formal description of chaining asymmetric capabilities) We focus on a key special case when the model is more accurate at verifying its own answers than it is at generating correct ones; that is, when the model exhibits a *verification-generation gap (VG Gap)*, on a particular problem domain [16]. In this section, we show that RL training on problem domains with VG gap (i) encourages chaining asymmetries, (ii) enables in-context exploration that (iii) discovers new solutions, often extrapolating to larger budgets and more difficult problem domains.

Definition 2.1 (Chaining asymmetric capabilities p, q in model π). Let $p, q : \mathcal{S} \mapsto \mathcal{S}$ be functions over token sequences \mathcal{S} (*e.g.*, p can be generation, q can be verification), and $\text{detect}(f, \tau)$ detects number of calls to function f in a token trace τ . For a reward r , we say that policy π chains asymmetries p, q if it benefits from calls to the composition $q(p(\cdot))$, compared to only $p(\cdot)$:

$$\mathbb{E}_{\tau \sim \pi} [r(\tau) \mid \text{detect}(q(p(\cdot)), \tau) > 0] > \mathbb{E}_{\tau \sim \pi} [r(\tau) \mid \text{detect}(p, \tau) > 0],$$

even though there is an optimal policy π_r^* that never calls q , *i.e.*, $\mathbb{E}_{\tau \sim \pi_r^*} \text{detect}[(q, \tau)] = 0$.

Setup. We validate the role of asymmetries in learning to explore by investigating two didactic tasks, on which Llama3.2-3B admits different VG gaps. First, the Countdown game [9, 8] (CDOWN) requires converting a set of numbers into an equation that outputs the desired target. The base LLM is more effective at verifying whether a proposed equation matches the target than searching over all possible equations to solve the task, and traces with more chained asymmetries are more performant, as we measure in Fig. 2, where performance on traces with more chains is higher. Second, we study n -digit multiplication MULT in natural language, without any external tools, where the base model exhibits limited verification (see App. E for asymmetry gap on MULT). Additionally, we supervise fine-tune Llama3.2-3B on correct n -digit multiplication traces containing verification to encourage more verification attempts (MULT-V); see App. E for details. Comparison of MULT vs. MULT-V enables direct evaluation of the benefits of asymmetries in base LLM, all else being held equal.

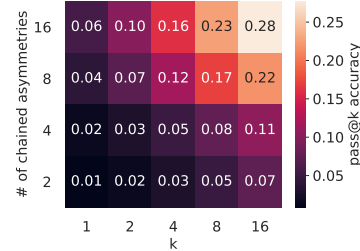


Figure 2: Measuring asymmetry (Definition 2.1) on CDOWN. Pass@k improves across k as the number of chained asymmetries increases.

1) Verification-generation asymmetry in π_b improves the performance of RL trained solutions. Fig. 3(a,b) shows a stark difference in performance and length of output traces as the training budget B_{tr} varies on CDOWN and MULT. On CDOWN, performance consistently increases as B_{tr} increases from 512 \rightarrow 2048, accompanied by a very clear increase in response length. On MULT, where the base model has limited propensity to verify, performance increases when B_{tr} increases from 1024 to 2048, but it plateaus thereon. Unlike CDOWN, test-time length is from saturating budget limits (and also oscillates in a widely RL training epochs). Contrast this with Fig. 3(d), where RL training on MULT-V, which leverages verification, exhibits longer lengths and stronger extrapolation performance. To summarize: **leveraging asymmetries improves performance and length-utilization in RL postraining.** Curiously, we also observe that models with greater VG gap exhibit less KL divergence from the base model; see App. E for figures and further discussion.

2) Chaining asymmetries enable extrapolation via in-context exploration. Interleaving verification and generation steps chains together asymmetric capabilities of the base model; we refer to this special case of skill-chaining as *chaining asymmetries*. To measure the benefits of chained asymmetries on CDOWN, we plot the pass@k accuracy of the base LLM, shown in Figure 2, and observe that performance increases with the chained asymmetries budget. Refer to App. E for more details. In Fig. 3 (c), we plot the extrapolation performance of the models trained at two values of B_{tr} . On CDOWN the model trained with B_{tr} 0.5-1k makes steady progress on problems in test budgets that are 8-16 \times B_{tr} itself. On MULT we find that B_{tr} has absolutely no effect on extrapolation performance with the base LLM that does not have VG asymmetry, but it has a substantial effect when the

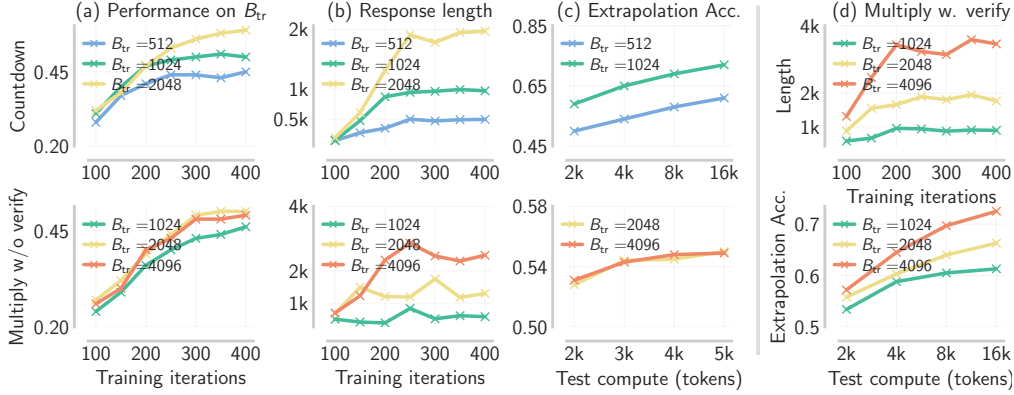


Figure 3: **RL training with and without asymmetries in π_b .** When asymmetries such as the VG gap are present (e.g., in CDOWN), RL training amplifies response length by chaining more asymmetries to explore in-context, where the probability of success improves with higher length on both B_{tr} and extrapolation regimes. On the other hand, when VG gap is absent in π_b (e.g., in MULT), increases in length and extrapolation performance are subdued. When we explicitly train on a base model fine-tuned to verify MULT (a setting we refer to as the MULT-V), we again observe upward length and extrapolation trends, consistent with CDOWN.

asymmetry is present. More importantly, while the base model without VG asymmetry fails to extrapolate and solve unsolved problems, with its accuracy improving by merely $\leq 2\%$ despite 16x compute scaling, the base model with VG asymmetry can still extrapolate well.

Why do asymmetries enable in-context exploration? We give an informal example of how an LLM can leverage VG gap to improve performance through longer in-context exploration: the p^k -model. We view the LLM as sequentially guessing candidate responses a_1, a_2, \dots , each with failure probability p , up to at most terminal k responses. We assume that this model admits perfect verification (*perfect VG gap*), which means that the learner can correctly assess whether each subsequent sequential response is correct and decide when to stop accordingly. In a simplified setting where attempts are independent, failure probability ($= p^k$) decays exponentially as k increases, as on CDOWN. In contrast, if verification is difficult (i.e., no VG gap), increasing k provides little benefit, since the model cannot adjudicate whether one guess is any better than another. In extreme scenarios, the only way to improve performance is by lowering p (better first guesses as seen on MULT).

Takeaways: Asymmetries are a critical pre-requisite for learning to explore.

- Asymmetries like the VG gap enable the model to continually explore, verify, and refine answers.
- RL training amplifies chaining of asymmetric skills and produces solutions that learn to explore in-context, thus benefiting from additional test-time compute beyond the training budget.

3 Negative Gradients Incentivize Exploration that Chains Asymmetries

Having observed that asymmetries in the base model are a prerequisite for in-context exploration, we ask: What enables models to exploit these asymmetries during RL? In this section, we show that a key ingredient in this is the *negative gradient*, the gradient term multiplied by a negative advantage (see Eq. 2 in App. B). Negative gradient drives in-context exploration via two mechanisms: (i) incentivizing the sampling of unseen token sequences; (ii) chaining asymmetries like VG gap (Sec. 2) that rapidly drives up response length and the number of verification attempts. For brevity, we refer to (i) as “exploration” [24] and (2) as in-context exploration (or “meta-exploration” [13, 25]).

Analysis setup. We analyze the evolution of response length, performance, and the number of chained asymmetries by comparing two training algorithms: (i) standard outcome-reward RL using GRPO [26] with token-level normalization [2]; (ii) GRPOMask, which zeros out (i.e. masks) the negative gradient and whilst retaining the *positive* gradient, thereby resembling online STaR [27] or RFT [28]. We conduct our experiments on CDOWN and DMATH reasoning (questions sourced from DeepScaleR dataset [29]) that exhibit the VG asymmetry. We make the following observations:

1) Negative gradient increases the number of chained asymmetries, and thereby boosts meta-exploration. When applied on an incorrect response y with tokens $y_1, y_2, \dots, \text{EOS}$, negative gradient reduces $p(y_i | y_{1:i-1})$, including $p(\text{EOS} | y)$ when the response ends before the training budget. Fig. 4(b) reveals that the probability mass recovered from the negative gradient (note: total probability is

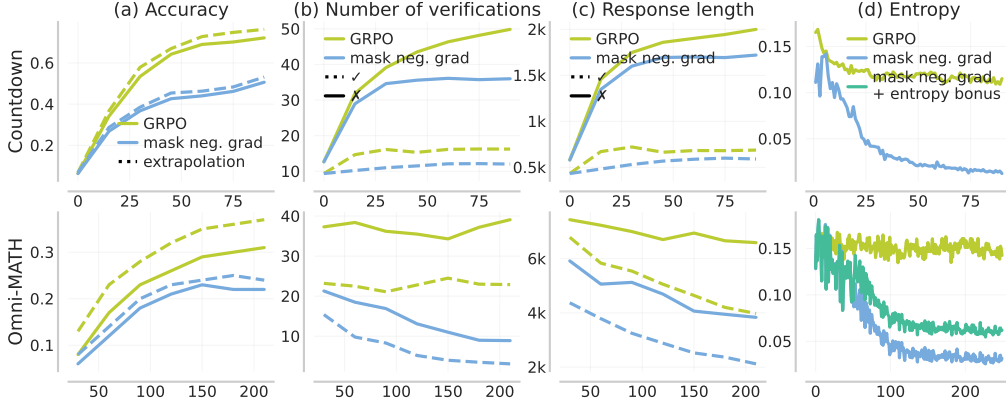


Figure 4: **RL training with and without negative gradients:** When the base model presents asymmetries, negative gradients promote in-context exploration by: (i) increasing length (shown in (c)) and chaining more asymmetries on incorrect responses, which shows up as more verification attempts (b); and (ii) increasing entropy and thus response diversity (d). This leads to better performance on the training budget and when extrapolating beyond it. In (b, c), ✓ denotes the statistic computed on correct responses and ✗ on incorrect responses.

conserved) is repurposed to increase the probability of chaining new pairs of asymmetric skills to the current trace (e.g., “Wait, ...” instead of terminating with EOS). When negative gradients are masked (GRPOMask) in CDOWN, we see that attempts (b) and length (c) plateau, accompanied by a decrease in performance. The relative trends between GRPO and GRPOMask are similar for DMATH, but differ in absolute terms. We include further discussion in App. F, where we also demonstrate that MULT (which does not exhibit asymmetries) benefits far less from negative gradients.

2) Negative gradients promote diverse responses during

RL training, encouraging exploration at two levels: (i) within a rollout; and (ii) across different rollouts. For (i), we observe that removing the negative gradient results in an entropy collapse over the next-token distribution (Fig. 4 (d)). This curtails diversity and leads to responses with a repeating stream of tokens when extrapolating the trained model to larger budgets (Fig. 5). For (ii), we measure the cumulative unique attempts on the CDOWN test set as we train the model (Fig. 5, see App. F for details) and find more unique attempts when training with negative gradients. While it is not surprising that RL algorithms benefit from exploration [30], it is unclear how this exploration can be particularly effective when extrapolating the model to larger test budgets.

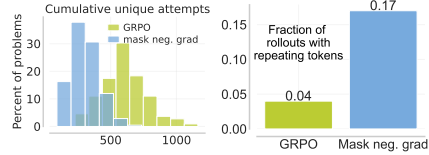


Figure 5: Negative gradients encourage distinct responses: they increase the cumulative number of unique attempts on CDOWN (left) and reduce responses that end with a repeating stream of tokens on DMATH (right).

3) LLMs trained with negative gradients extrapolate better. The bridge between exploration and meta-exploration lies in the use of asymmetries present in the base model. Exploration afforded by negative gradients, in the presence of asymmetries like the VG gap incentivizes meta-exploration, because longer responses with more chained asymmetries (verification-generation steps) discover correct solutions and get positively rewarded. Recall from Sec. 1 that if a model has indeed learned to explore in-context or meta-exploration, it should benefit from additional test compute since more search should only improve performance under high VG gaps. We confirm this in Fig. 4(a) (dotted lines), where we see that when testing on hard test problems in DMATH, on a budget that is $2 \times B_{tr}$, the performance gap widens with negative gradients, in comparison to the masked version. We include an additional analysis of negative gradient dynamics in a didactic setting (p^k -model mentioned at the end of the previous section) in App. C.

Takeaways: Negative gradient incentivizes in-context exploration under large VG gaps

- Negative gradients in RL “move” probability from early stopped incorrect answers onto other likely modes, e.g., those that exploit asymmetries in the base LLM. When the VG gap is large, negative gradients force additional attempts and longer responses that eventually discover the correct answer. As a result, in-context exploration is positively rewarded and reinforced.
- Negative gradients boost response diversity and thus coverage over correct answers, as confirmed by our empirical results on CDOWN, DMATH, and theoretical results in the bi-gram model.

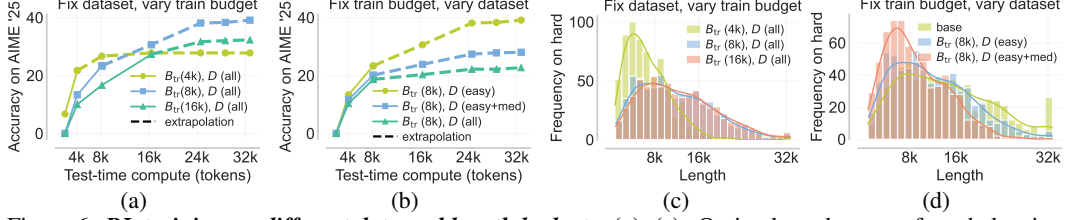


Figure 6: **RL training on different data and length budgets.** (a), (c): Optimal results come from balancing optimization difficulty (better at shorter budgets) and in-context exploration (better at longer budgets). (b), (d): Training on hard problems at the 8k budget kills in-context exploration required by hard problems

4 Coupled Curriculum Training Structures Exploration in Long Length RL

In the presence of asymmetries, training with negative gradients produces models that can extrapolate beyond their training budget. However, just negative gradients are not enough: as we see in Fig. 6(a), different training budgets B_{tr} lead to different levels of performance on B_{tr} , as well as extrapolated test compute. *So how should we set the budget B_{tr} to attain strong extrapolation performance?* And in correspondence with token budgets, *what prompts should we be training on for a given budget?*

Setup. We evaluate on DMATH, CDOWN, with different training budgets and data compositions. We split DMATHEvenly across three levels of hardness by Qwen-R1-Distilled-32B accuracy. For CDOWN, we judge problem difficulty based on the number of terms in the equation. We use the GRPO [26] algorithm to train models on all compute budgets and datasets (see App. G for details).

Training solely at low or high values of B_{tr} is not desirable. We first run RL training for 300 iterations on the easy DMATH problems at different training budgets $B_{tr}=4k, 8k, 16k$ (see Fig. 6(a)). While training at the short budget $B_{tr}=4k$ attains the best performance at the same test budget of 4k tokens, it “kills” exploration and leads to poor extrapolation performance (no gains from 8k to 40k). This is because the budget is saturated too quickly, so potentially correct responses get negatively rewarded, hindering length increase and asymmetry-chaining driven by the negative gradient. Fig. 6(c) reveals that this biases the model to stop early and terminate incorrectly, failing to exploit the VG gap. On the other extreme, training at $B_{tr}=16k$ introduces significant optimization challenges, typical of long horizons policy gradients suffering from high gradient variance [21]. This model performs worse on its own training budget of 16k than it does training on $B_{tr}=8k$ and extrapolating to 16k. We find that $B_{tr}=8k$ attains the best scaling when extrapolating test compute, implying that we need to strike a balance between the length budget available for negative gradient to encourage chaining asymmetries (infeasible in $\leq 4k$ tokens) and mitigating optimization challenges.

Training naively on a static data mixture is insufficient. Having identified a reasonable B_{tr} of 8k, we now turn to studying the effect of data compositions. We compare the naïve training data mixture with all difficulties (easy+med+hard) against easy, easy+med at $B_{tr}=8k$. Matching train and test composition is ideal for better “in-distribution” performance, *i.e.*, when evaluating models at B_{tr} (see App. G). Surprisingly, the same is not true for extrapolation on out-of-distribution (OOD) problems at larger test-time budgets. As shown in 6(b), the model trained on *only easy* problems obtains the best performance on OOD AIME’25 when extrapolating compute to 40k. This is despite the fact that AIME’25 problems resemble harder problems and (few prior AIME problems are also present in DMATH). Why does this happen? Given a dataset, training on budgets smaller than the length of a typical base model response on that dataset penalizes in-context exploration early in training, resulting in overly short solutions that do not learn to chain asymmetries and explore. (Fig. 6(d)).

How can we avoid challenges with training on a fixed dataset and length budget? One approach is to incorporate a budget curriculum that varies B_{tr} over training. However, this alone is insufficient because, as shown above, training on hard problems with short budgets suppresses length and in-context exploration. On the other hand, we can design a curriculum over the difficulty level and keep the training budget fixed at a high enough value. However, this leads to learning over-exploratory traces tailored to easy problems (see App. G for a detailed study of this on CDOWN). We also show an experiment comparing our proposed fix below with only budget or data curricula in Fig. 7(d).

Our recipe e3: coupled curriculum for test-time extrapolation. Motivated by our findings above, we propose a **coupled curriculum** which varies B_{tr} and problem difficulty in a coordinated fashion as training progresses. We refer to our entire recipe as e3: *exploration enables extrapolation*. This

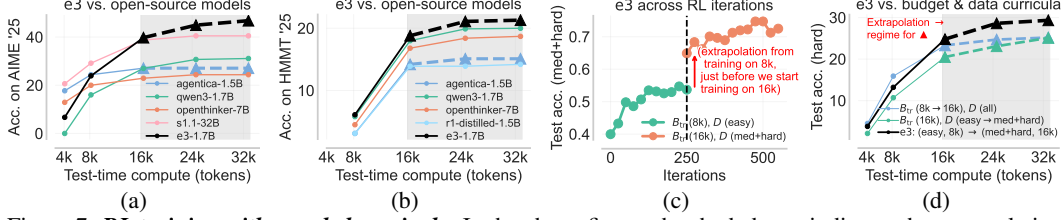


Figure 7: **RL training with coupled curricula.** In the above figure, the shaded area indicates the extrapolation regime. (a), (b): e3 achieves state-of-the-art performance across models < 2B. (c): extrapolation gain from switching to a longer budget during training (d): coupled curriculum outperforms data and budget curricula

encompasses both the description of best-practices that enable in-context exploration — asymmetries and negative gradient — with a prescription for the coupled curriculum we describe next.

Key principle for curriculum design. We simplify curriculum design by first fixing the dataset at each stage: moving from easy to hard. Now, the key question is to select the appropriate budget for dataset D_i at stage i of the curriculum. Intuitively, we want a budget $B_{tr,i}$ such that training on $B_{tr,i}$ positively rewards in-context exploration, which will in turn improve extrapolation performance and provide a good initialization for the next stage $i + 1$. At the same time, to be optimization friendly, we want $B_{tr,i}$ to be minimal while being large enough to accommodate most responses from the given model π_i . Encoding these conditions on $B_{tr,i}$, we propose the following optimization problem.

$$B_{tr,i}^*(D_i) = \arg \min_B B \text{ s.t. } J(\pi_i; D_i, 2 \cdot B) \geq \kappa \cdot J(\pi_i; D_i, B) \text{ and } B \geq \mathbb{E}_{\mathbf{x} \sim D_i, \tau \sim \pi_i(\cdot|\mathbf{x})} |\tau|, \quad (1)$$

where $J(\pi_b; D, B_{tr})$ denotes the performance of the base model π_b at budget B_{tr} on dataset D , and $|\tau|$ denotes token count. In practice, we solve the above problem over a fixed set of training budgets: 4k, 8k, 16k, and find this to be a useful heuristic to greedily choose $B_{tr,i}$ in a way that incentivizes in-context exploration at stage i , since it is hard to jointly optimize the budgets across all stages. *E.g.*, setting $\kappa = 1.2$, we find 8k to be the optimal choice for training on easy problems (note that the trained model also satisfies the condition in Eq. 1 at $\kappa = 1.2$, see Fig. 6(a)). Following this, our recipe e3 fine-tunes the base model on easy problems in DMATH at a training budget of 8k, and subsequently continues training on medium and hard problems in DMATH with a budget of 16k.

Final results with e3. In Fig. 7(a,b), we compare the performance of a Qwen3-1.7B model fine-tuned using e3 with open-source models, including some 7B and 32B models. As shown, e3 achieves state-of-the-art performance on AIME'25 and HMMT'25, within a model class of size < 2B. We outperform the best model in this class by > 10% on AIME'25 in terms of peak performance, and show that our model, trained only up to a budget of 16k, extrapolates better than other models including s1.1-32B [4] and OpenThinker-7B [3] when we extend the compute budget up to 32k. In Fig. 7(c), we show that the model already learns to extrapolate at a point during training when we move from the 8k budget to the 16k budget, where there is a > 10% performance gain. Finally, Fig. 8 shows that compared to budget forcing, which is a prompting technique introduced in s1 [4] to enable extrapolation, e3 achieves significantly better scaling, even without applying budget forcing to it. Refer to Appendix L for a discussion on the conclusion and limitations.

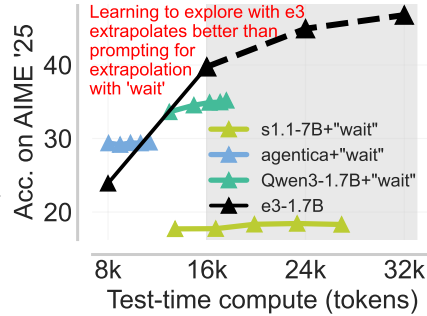


Figure 8: e3 (w/o “wait”) is superior when extrapolating to larger budgets, compared to budget forcing with “wait” prompt 2/4/6/8 times.

Takeaways: Coupled data & budget curriculum structures exploration during RL.

- RL with fixed B_{tr} , D hurts in-context exploration: (i) short B_{tr} penalizes exploration on hard problems as budget is overrun; (ii) large B_{tr} overfits to under-exploratory behavior on easy ones.
- We propose a coupled curriculum e3: at each stage, given D , choose smallest B_{tr} such that chaining more asymmetries till a budget of $2 \cdot B_{tr}$ is positively rewarded at RL initialization.
- By fine-tuning Qwen3-1.7B with e3, we outperform any < 2B model on AIME'25, HMMT'25.

References

- [1] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>.
- [2] Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
- [3] OpenThoughts Team. Open Thoughts. <https://open-thoughts.ai>, February 2025.
- [4] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling, 2025. URL <https://arxiv.org/abs/2501.19393>.
- [5] Tianle Li, Ge Zhang, Quy Duc Do, Xiang Yue, and Wenhui Chen. Long-context llms struggle with long in-context learning. *arXiv preprint arXiv:2404.02060*, 2024.
- [6] Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via reinforcement learning, 2024. URL <https://arxiv.org/abs/2409.12917>, 2024.
- [7] Amrith Setlur, Yuxiao Qu, Matthew Yang, Lunjun Zhang, Virginia Smith, and Aviral Kumar. Optimizing llm test-time compute involves solving a meta-rl problem. <https://blog.ml.cmu.edu/>, 2025. CMU MLD Blog.
- [8] Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and Noah D Goodman. Stream of search (sos): Learning to search in language. *arXiv preprint arXiv:2404.03683*, 2024.
- [9] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023.
- [10] Yuxiao Qu, Matthew YR Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan Salakhutdinov, and Aviral Kumar. Optimizing test-time compute via meta reinforcement fine-tuning. *arXiv preprint arXiv:2503.07572*, 2025.
- [11] Andreas Hochlehnert, Hardik Bhatnagar, Vishaal Udandaraao, Samuel Albanie, Ameya Prabhu, and Matthias Bethge. A sober look at progress in language model reasoning: Pitfalls and paths to reproducibility. *arXiv preprint arXiv:2504.07086*, 2025.
- [12] Dibya Ghosh, Jad Rahme, Aviral Kumar, Amy Zhang, Ryan P. Adams, and Sergey Levine. Why generalization in RL is difficult: Epistemic pomdps and implicit partial observability. *CoRR*, abs/2107.06277, 2021. URL <https://arxiv.org/abs/2107.06277>.
- [13] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [14] Amrith Setlur, Nived Rajaraman, Sergey Levine, and Aviral Kumar. Scaling test-time compute without verification or rl is suboptimal. *arXiv preprint arXiv:2502.12118*, 2025.
- [15] Gokul Swamy, Sanjiban Choudhury, Wen Sun, Zhiwei Steven Wu, and J Andrew Bagnell. All roads lead to likelihood: The value of reinforcement learning in fine-tuning. *arXiv preprint arXiv:2503.01067*, 2025.
- [16] Yuda Song, Hanlin Zhang, Carson Eisenach, Sham Kakade, Dean Foster, and Udaya Ghai. Mind the gap: Examining the self-improvement capabilities of large language models. *arXiv preprint arXiv:2412.02674*, 2024.
- [17] Seungone Kim, Ian Wu, Jinu Lee, Xiang Yue, Seongyun Lee, Mingyeong Moon, Kiril Gash-teovski, Carolin Lawrence, Julia Hockenmaier, Graham Neubig, et al. Scaling evaluation-time compute with reasoning models as process evaluators. *arXiv preprint arXiv:2503.19877*, 2025.

- [18] Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D. Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars, 2025. URL <https://arxiv.org/abs/2503.01307>.
- [19] Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, and Aviral Kumar. Preference Fine-Tuning of LLMs Should Leverage Suboptimal, On-Policy Data, ICML 2024.
- [20] Yi Ren and Danica J Sutherland. Learning dynamics of llm finetuning. *arXiv preprint arXiv:2407.10490*, 2024.
- [21] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76, 2021.
- [22] Allen Nie, Yi Su, Bo Chang, Jonathan N Lee, Ed H Chi, Quoc V Le, and Minmin Chen. Evolve: Evaluating and optimizing llms for exploration. *arXiv preprint arXiv:2410.06238*, 2024.
- [23] Akshay Krishnamurthy, Keegan Harris, Dylan J Foster, Cyril Zhang, and Aleksandrs Slivkins. Can large language models explore in-context? *arXiv preprint arXiv:2403.15371*, 2024.
- [24] Susan Amin, Maziar Gomrokchi, Harsh Satija, Herke Van Hoof, and Doina Precup. A survey of exploration methods in reinforcement learning. *arXiv preprint arXiv:2109.00157*, 2021.
- [25] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, and S. Levine. Meta-reinforcement learning of structured exploration strategies. *CoRR*, abs/1802.07245, 2018.
- [26] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- [27] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.
- [28] Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. Scaling relationship on learning mathematical reasoning with large language models. *arXiv preprint arXiv:2308.01825*, 2023.
- [29] Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Tianjun Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL. <https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2>, 2025. Notion Blog.
- [30] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. In *ICML*, 2019. URL <https://arxiv.org/pdf/1812.02690.pdf>.
- [31] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- [32] Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang Geng, Jacob Eisenstein, Rishabh Agarwal, Alekh Agarwal, Jonathan Berant, and Aviral Kumar. Rewarding progress: Scaling automated process verifiers for llm reasoning. *arXiv preprint arXiv:2410.08146*, 2024.
- [33] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
- [34] Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching language model agents how to self-improve. *arXiv preprint arXiv:2407.18219*, 2024.
- [35] Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.

- [36] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, and Zonghan Yang. Kimi k1.5: Scaling reinforcement learning with llms, 2025. URL <https://arxiv.org/abs/2501.12599>.
- [37] OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. Openai o1 system card, 2024. URL <https://arxiv.org/abs/2412.16720>.
- [38] Hugging Face. Open r1: A fully open reproduction of deepseek-r1, January 2025. URL <https://github.com/huggingface/open-r1>.
- [39] 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>.
- [40] Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild, 2025. URL <https://arxiv.org/abs/2503.18892>.
- [41] Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv preprint arXiv:2503.18892*, 2025.
- [42] Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A fully open-source 14b coder at o3-mini level. <https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349a512a51>, 2025. Notion Blog.
- [43] Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with reinforcement learning. *arXiv preprint arXiv:2503.04697*, 2025.
- [44] Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.
- [45] Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, and Alekh Agarwal. A minimaximalist approach to reinforcement learning from human feedback. *arXiv preprint arXiv:2401.04056*, 2024.
- [46] Qiyang Li, Yuexiang Zhai, Yi Ma, and Sergey Levine. Understanding the complexity gains of single-task rl with a curriculum. *arXiv preprint arXiv:2212.12809*, 2022.
- [47] Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Lucas Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, et al. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*, 2025.
- [48] Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2502.14768*, 2025.
- [49] Taiwei Shi, Yiyang Wu, Linxin Song, Tianyi Zhou, and Jieyu Zhao. Efficient reinforcement finetuning via adaptive curriculum learning, 2025. URL <https://arxiv.org/abs/2504.05520>.
- [50] Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acemath: Advancing frontier math reasoning with post-training and reward modeling. *arXiv preprint*, 2024.
- [51] Arash Ahmadian, Chris Cremer, Matthias Gall  , Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet   st  n, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms. *arXiv preprint arXiv:2402.14740*, 2024.

- [52] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [53] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv: 2409.19256*, 2024.
- [54] Yang Gao, Christian M Meyer, Mohsen Mesgar, and Iryna Gurevych. Reward learning for efficient reinforcement learning in extractive document summarisation. *arXiv preprint arXiv:1907.12894*, 2019.

Appendices

- A.** Related work.
- B.** Optimizing & Extrapolating Test-Time Compute.
- C.** Analyzing Negative Gradient Dynamics in the p^k Model.
- D.** Testing Extrapolation of Open Source Models.
- E.** Additional Experiments and Details for Section 2 (Chained Asymmetries).
- F.** Additional Experiments and Details for Section 3 (Negative Gradient).
- G.** Additional Experiments and Details for Section 4 (Curricula Training).
- H.** Omitted Proofs.
- I.** Broader Impact Statement.
- J.** Note on Computational Resources Used for e3.
- K.** Example Traces.
- L.** Conclusion and Limitations.

A Related Work

Scaling test-time compute via Long CoT reasoning. Prior work explores a number of avenues for scaling test-time compute, including majority voting [31], best-of-n sampling, and beam search [32, 33], as well as sequential self-correction [34, 35]. More recent results indicate that training models to use test-time compute to generate longer chains of thought (CoT) that combine verification, search, and self-correction – all in a free-form manner, performs better [1, 36, 37], resulting in widespread open-source reproduction efforts [38–40, 29]. We situate our work in the paradigm of long CoT reasoning and study the role of algorithms (RL or SFT), data composition, and design of the training procedure.

Test-time extrapolation. The true benefit of test-time scaling is consistently improving performance as we extrapolate test compute. While prior work tests the model’s performance on budgets longer than the training budget [41, 42], they do not explain the relationship between the training recipe and the extrapolation capabilities. In our work, we provide a clear recipe and explain the mechanism behind why our recipe enables test time extrapolation. Other works perform extrapolation by explicitly prompting models to generate more tokens when a response terminates [4, 43]. In this work, we show that models that learn to explore in-context extrapolate test compute better than prompting-based approaches. In particular, we study the role played by the base model, training algorithm (RL), as well as data mixtures and token budgets, on the ability to extrapolate. Furthermore, prior work [14] has investigated how performance scales with budgets when train and test budgets are matched, which is different from the OOD setting this work where test budgets are significantly longer.

Exploration in test-time scaling. Long CoTs allow models to explore various reasoning paths before exploiting and committing to a final answer. While prior works have shown the importance of the base model’s ability to conduct exploration [18, 44], we discover the crucial enabling factor is the presence of *asymmetries* in the model. Next, we show that the negative gradient in RL incentivizes the model to chain together multiple asymmetries, which in turn leads to an increase in the length of the response. In contrast, SFT alone does not provide this kind of chaining or exploration benefits. Our analysis is orthogonal to theoretical works Setlur et al. [14], Swamy et al. [45], which shows that RL performs better than SFT, but from a statistical perspective, whereas our argument is more focused on the learning dynamics. Concurrent work builds techniques to boost exploration during RL via advantage normalization [46, 2] or PPO clipping [2], and these techniques can be combined with e3, but they do not highlight the role of negative gradients in learning to explore. Finally, Wang et al. [47] briefly remarks about the role of policy gradient loss and entropy when running RL with only a few examples. Our study formally investigates the underlying mechanism of negative gradients increasing length and entropy with controlled experiments and theoretical results.

Data and length curricula. Recent works have also investigated using a curriculum on problem difficulty [36, 48, 49] and context window length [29, 50] during RL training. Their motivation stems primarily from an efficiency standpoint: avoiding zero advantage updates [49, 2], efficient optimization [29], or efficiency of using test-time compute [10]. While we do make similar observations regarding each curriculum individually, perhaps our most interesting finding is that carefully coupling both data and budget curricula can lead to much better performance and extrapolation, beyond merely some gains in compute efficiency. We show that training on hard problems with

short budgets often yields terse solutions that fail to extrapolate, while easy problems with long budgets can cause optimization issues or verbose outputs. Thus, curricula must be carefully designed to support effective extrapolation. Conceptually, our curricula are most related to dense progress rewards [10, 32], in the sense that curricula incentivize different degrees of progress for different questions, at different points in training. We believe this is a good avenue for future work to pursue.

B Optimizing & Extrapolating Test-Time Compute

Post-training scaling test-time compute. RL and SFT are categories of post-training algorithms that refine a pre-trained base LLM π_b into a reasoning model, in particular one that utilizes more test-time compute by producing long chains-of-thought to succeed. Typical outcome-reward RL trains LLM π (initialized with π_b) to maximize performance on outcome 0/1 reward $r^*(\mathbf{x}, \mathbf{y})$, for inputs $\mathbf{x} \sim \rho$ and response $\mathbf{y} \sim \pi(\mathbf{y} | \mathbf{x})$ restricted to an apriori fixed maximum token length or *training budget* B_{tr} [2, 29]. On the other hand, SFT fine-tunes π_b on long thinking traces from more capable models or humans to distill their reasoning capabilities [3, 4], where the maximum length of the expert traces also implicitly induces a training budget B_{tr} , similar to RL. **Our goal**, is to train models that can improve performance when we extrapolate test-compute beyond B_{tr} . Even though the true promise of test-time compute is extrapolation performance, we find that *current thinking models fall short on extrapolation*.

We evaluate multiple models on a test budget of 32K, $\approx 1.5\text{--}2 \times B_{tr}$ across all models. We plot our results on AIME25 in Fig. 9 (see App. D for a detailed comparison) and note that most of the performance gains lie in the training budget, and the gains are minuscule as we test beyond that.

Negative gradient in RL. A key distinction between SFT and RL is the *negative gradient*, which corresponds to the part of the policy gradient coming from traces that fail. In Eq. 2 we present a generalized version of the policy gradient adopted by most RL post-training methods: REINFORCE [51], PPO [52], and GRPO [26]. From this we note that on a prompt \mathbf{x} , RL training observes two types of gradients: (i) the positive gradient which maximizes the likelihood of a correct responses \mathbf{y} with a positive advantage $A(\mathbf{x}, \mathbf{y})$, and (ii) the negative gradient which *pushes down* the likelihood of an incorrect response with a negative advantage $A(\mathbf{x}, \mathbf{y})$. Here, \mathbf{y} can be sampled *on-policy* $\pi = \tilde{\pi}$ or *off-policy* $\pi \neq \tilde{\pi}$. Thus, we can view SFT as a purely positive gradient method that only maximizes likelihood on correct reasoning traces. In Sec. 3, we show why the negative gradient is largely responsible for driving up response lengths and in-context exploration during RL, thereby enabling RL-trained models to explore more at test-time and extrapolate better compared to SFT-based ones.

$$\mathbb{E}_{\mathbf{y} \sim \tilde{\pi}(\cdot | \mathbf{x})} [A_i(\mathbf{x}, \mathbf{y}) \cdot \nabla_{\pi} \log \pi(\mathbf{y} | \mathbf{x})] \quad (\text{general form of policy gradient in RL}) \quad (2)$$

C Analyzing Negative Gradient Dynamics in the p^k Model.

Recall from Sec. 2 the p^k model, where an LLM samples k independent actions sequentially, verifies them (with a perfect accuracy), and terminates immediately after the correct one is produced. In this section, we introduce a didactic setup where verification is perfect but attempts are not independent, akin to LLMs we train in practice.

We consider a Markov decision process (MDP) with action space $\bar{\mathcal{A}} = \mathcal{A} \cup \{\text{stop}\}$, where $\mathcal{A} = [100]$ are standard actions and *stop* is an early “stopping” action (like EOS) that terminates the trace. For simplicity, we consider policies parametrized as a softmax bigram model $\pi_M(a_{t+1} | a_t)$, with details deferred to the App. F. In this bi-gram model, the current state s_t always matches the previous action a_{t-1} , and $a^* \in \mathcal{A}$ denotes the optimal action. In a rollout a_1, \dots, a_t , the initial action a_1 is sampled from a fixed π_0 . For $t > 1$, a learner policy samples an action $a_t \sim \pi(\cdot | a_{1:t-1}) \in \Delta(\mathcal{A})$. The MDP terminates with reward 1 at time t if $a_t = a^*$, and with reward 0 if $a_t = \text{stop}$ (stops too early), or $t > B_{tr}$ (budget is exhausted before a correct response). The model **learns to explore** if it learns to never play *stop* for any t (no early stopping), until a^* is observed, *i.e.*, increasing k in p^k . **Refining the guess** amounts to upweighting $\pi(a^* | a_{1:t-1})$ without reducing $p(\text{stop})$, *i.e.*, improving p in p^k .

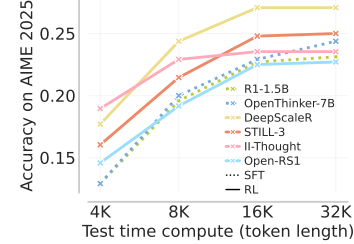


Figure 9: Accuracy on AIME 2025 of various open-source models at different test time compute budgets. Performance gains diminish as the test-time budget increases, with virtually no gains from 16k to 32k.

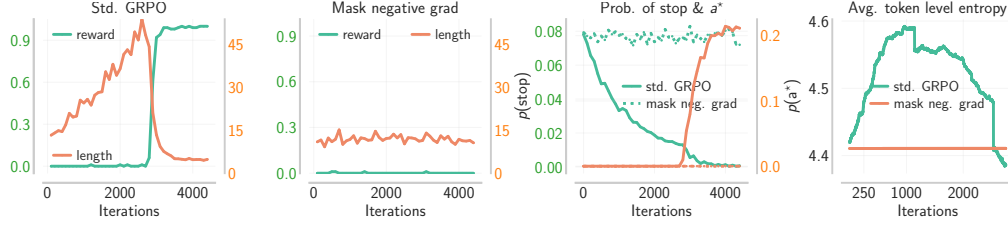


Figure 10: *Negative gradients in a bi-gram model.* Negative gradients push down $p(\text{stop})$ during training (c), increasing length (a) and entropy of the next action distribution (d) to accommodate more in-context exploration, only decreasing them when a^* is discovered. In contrast, positive gradients rarely change $p(\text{stop})$ or entropy.

1) Negative gradient increases length until $p(a^*)$ is reasonably high. In Fig. 10(a) standard GRPO ($B_{\text{tr}}=100$) increases average response length from 15 to 45 at budget, driven by the drop in the marginal probability of stopping early $p(\text{stop})$ (Fig. 10(c)). After multiple RL iterations with negative gradients, the average number of attempts per trace is sufficiently large, and the learner can sample a^* with non-trivial probability in any given trace. Once this happens (Fig. 10(c)), in our simple bigram setup, the model rapidly upweights the likelihood of one-step transitions to a^* , resulting in a phase transition where reward increases as length drops. In our LLM benchmarks, however, we do not see the same phase transition since finding “shortcuts” to correct responses is considerably more difficult. In contrast, GRPOMask (Fig 10(b)) fails to improve reward or increase length.

2) Negative gradient improves coverage by increasing entropy of $\pi(\cdot \mid a_{1:t-1})$. When π_M samples a highly likely yet incorrect action, the negative gradient computed on this sample increases entropy by moving probability mass onto less-seen modes of the distribution, including a^* . We show this formally in Theorem C.1 where we prove that upon sampling a highly likely incorrect action with probability p , GRPO update with a negative gradient results in an entropy increase of $\approx p^2$ when all other actions, including a^* are highly unlikely. We note this empirically as well in Fig. 10(d), where conditional entropy increases across states, until a^* is discovered, after which it drops sharply as the positive gradient rapidly moves mass onto a^* within a few iterations.

Theorem C.1 (Negative gradient increases entropy when a^* is unlikely; formal version of Thm. H.3). *At state s , if the most likely action under π is $a_1 =: \arg \max_{a'} \pi(a'|s) \neq a^*$, then, for any π , a negative stochastic gradient step increases the entropy of $\pi(\cdot|s)$ with prob. $\geq \pi(a_1|s)$. Additionally, in a suitable regime of π , the increase $\gtrsim (\pi(a_1|s) - \pi(a_2|s))^2$, where a_2 is second most likely after a_1 . In contrast, in the absence of the negative gradient, the entropy is preserved with prob. $1 - \pi(a^*|s)$.*

D Testing Extrapolation of Open-Source Models

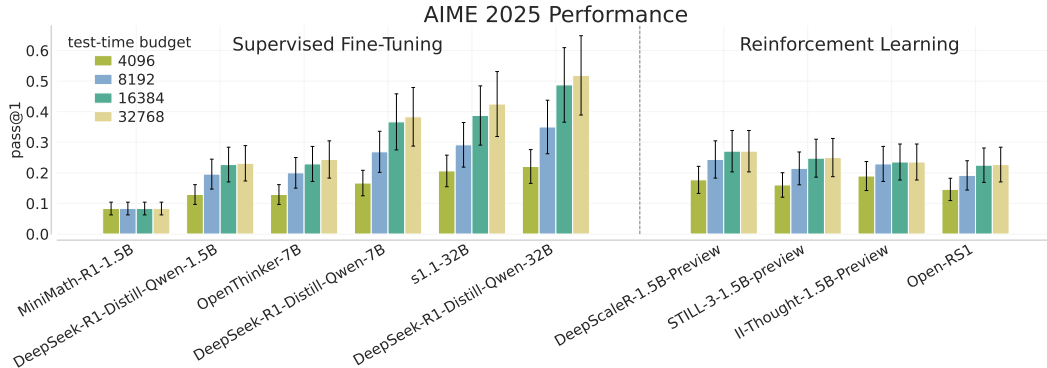


Figure 11: *Extrapolation of test-time compute:* We plot the performance (pass@1) on AIME 2025 at different test-time compute budgets across multiple open-source models of different sizes, trained with SFT or RL.

Extrapolation on AIME 2025 Extrapolation (i.e. the chaining of generation, verification, refinement, etc.) can potentially extend LLM performance after training, and do so beyond the context length the model was originally trained on. To evaluate this properly, we need sufficiently challenging problems that allow meaningful expressiveness in reasoning beyond small context lengths. The math

problems associated with AIME align with this, and our evaluations prioritize AIME 2025 to attempt to mitigate any potential data contamination in the models’ training sets from previous years of AIME. The goal of the experiment is to measure the extent to which test-time compute influences overall model performance as context length increases, with the expectation that increasing output length allows models to "reason" for longer periods, continuing the extrapolation process, and ultimately arriving at the correct answer more frequently.

Experiment Setup Inference for every open-source model was performed using Oumi through data-parallel SGLang. All models had inference run with a max output length of approximately 32k tokens, though some are slightly lower due to this exceeding their max context length when combined with the prompt. The exact inference hyperparameters are described in Table 1. After inference, the model responses were truncated from the right side until the number of remaining tokens present was equal to the specified test-time budget. 16 responses were collected for every problem in AIME with the specified inference settings, and the Pass@1 rate was calculated by averaging over these 16 responses. Final answers were extracted using a regular expression for the boxed portion of the answer, with correct answers marked as passing and incorrect or incorrectly parsed answers marked as nonpassing. The prompt used is in Box D.1, and the problems were taken from the FVU AIME 2025 dataset on HuggingFace¹.

Box D.1: AIME Evaluation Prompt Template

You will be given a math problem. Solve the problem step by step. Output your final answer in the form of `\boxed{your answer}`. Problem: {problem}

Model	Temp.	Top p	Rollouts	Max New Tokens	Model Max Length
MiniMath R1-1.5B	0.6	0.95	16	32768	40960
DeepSeek R1-Distill-Qwen-1.5B	0.6	0.95	16	32768	40960
OpenThinker-7B	0.6	0.95	16	31000	32768
DeepSeek-R1-Distill-Qwen-7B	0.6	0.95	16	32768	40960
s1.1-32B	0.6	0.95	16	31000	32768
DeepSeek-R1-Distill-Qwen-32B	0.6	0.95	16	32768	40960
DeepScaleR-1.5B-Preview	0.6	0.95	16	32768	40960
STILL-3-1.5B-preview	0.6	0.95	16	32768	40960
II-Thought-1.5B-Preview	0.6	0.95	16	32768	40960
Open-RS1	0.6	0.95	16	32768	40960

Table 1: Inference parameters used for generating the extrapolation plots in Figure 9.

Results The results in Figure 11 show that as the maximum number of output tokens increases, every model capable of "reasoning" is able to attain a higher Pass@1 rate, with performance generally saturating at 16k tokens with relatively minor improvements at 32k. We do not observe this with MiniMath-R1-1.5B, and we suspect this is due to its fine-tuning focusing solely on smaller math problems trained with supervised fine-tuning, likely resulting in catastrophic forgetting of the ability to continuously extrapolate. Interestingly, we do not see a strong improvement in extrapolation behavior among models tuned with reinforcement learning compared to DeepSeek R1-Distill-Qwen-1.5B, which was trained with supervised fine-tuning. We suspect that this is likely due to the nature of the distillation data from the R1 model, which, if varied sufficiently in length, could avoid the length bias normally learned from supervised fine-tuning, while still teaching the model to perform extrapolation.

E Additional Experiments and Details for Section 2 (Chained Asymmetries)

E.1 Details on MULT and MULT-V

Data collection. Both MULT and MULT-V consist of multiplication traces for solving a 5-digit \times 5-digit multiplication problem. For the MULT task, we use a Llama3.2-3B instruction tuned model

¹https://huggingface.co/datasets/FVU/AIME_2025

where the number of intermediate verification attempts is much lower in a trace when asked to solve a multiplication problem. In fact, it is not hard to see that, in general, for multiplication, generation of a trace may be as hard as verifying a generated one, as the only way to verify the entire trace is to re-attempt the multiplication or carry out a division with the computed target. We contrast this task with the MULT-V task, where the Llama3.2-3B models are first finetuned on traces from Qwen-32B-R1-Distilled and GPT-4o models. These traces contain multiple verification attempts that verify intermediate steps solving smaller multiplication problems, and the steps are part of an entire trace that attempts to solve the main multiplication problem involving two 5-digit numbers. For collecting data we used the prompt in Box E.1. In App. K Example 2, we also provide an example multiplication trace with verification attempts sampled by the base model in MULT-V. As we will see in Fig. 15, the absence of asymmetries in MULT leads to lower accuracy and verifications when compared to MULT-V, where asymmetries are present.

Box E.1: Prompt for generating MULT-V data

Multiply {num1} and {num2}. Please reason step by step, and put your final answer within $\boxed{}$. At each step, try to verify your response if possible and prefix the line with "Check:". <think>

Hyperparameter	Values
train_batch_size	256
ppo_mini_batch_size	64
learning_rate	5.0e-6
kl_loss_coef	0.001
entropy_coeff	0.001
temperature	1.0
rollout.n	16

Table 2: Verl [53] hyperparameters used for MULT and MULT-V.

Training details. Hyperparameters for our experiments on MULT and MULT-V are given in Table 2.

E.2 Details on CDOWN

Training details. Hyperparameters in CDOWN experiments follow the table below unless otherwise specified. In all of our CDOWN experiments, we take the fine-tuned Llama3.2-3B base model from [18]. For Fig. 3, we trained with $B_{tr} = 512, 1024, 2048$ on problems with 3, 4, 5, 6 candidates. The total number of datapoints we used was 40000, which were evenly split across the four difficulties.

Hyperparameter	Values
train_batch_size	128
ppo_mini_batch_size	32
learning_rate	1.0e-6
kl_loss_coef	0.001
entropy_coeff	0
temperature	0.6
rollout.n	8

Table 3: Verl [53] hyperparameters used for CDOWN.

Evolution of chained asymmetries at test time. In Fig. 12, we show that as training progresses, responses with more chained asymmetries enjoy a greater improvement. If we move across any diagonal parallel to the main diagonal from top left to bottom right, we move across a constant attempt budget (e.g., moving from 16 chained asymmetries \times 1 pass to 8 chained asymmetries \times 2 passes). Having sequential chained asymmetries become increasingly better than parallel rollouts as training progresses, indicating the exploitation of asymmetries in RL training. See example of chained asymmetry in App. K, Example 1.

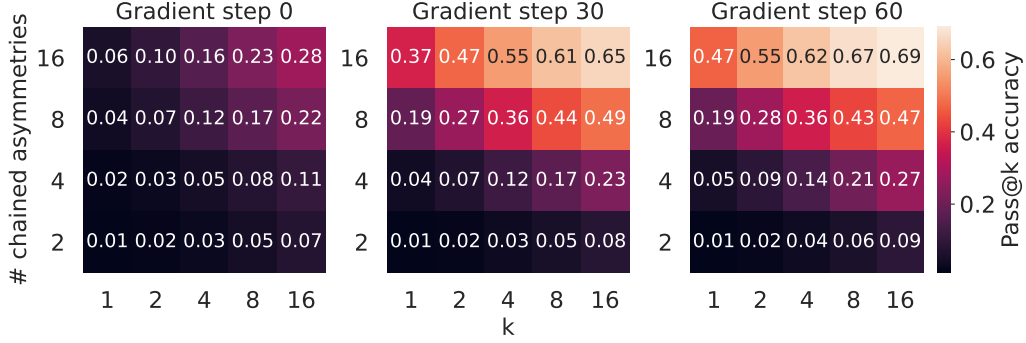


Figure 12: *Evolution of asymmetries during training on CDOWN*: More chained asymmetries lead to a greater improvement in pass@k performance across gradient steps.

E.3 In the Presence of Asymmetries, KL Divergence with Base LLM Reduces as Training Budget Increases

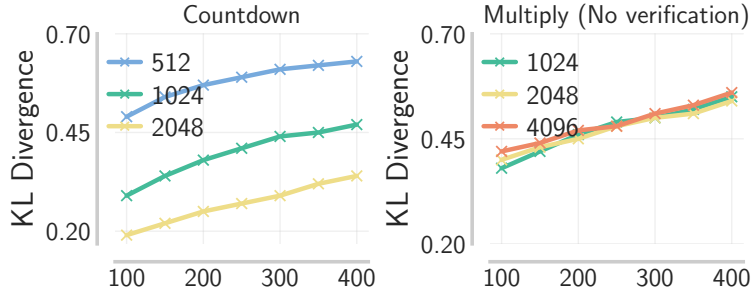


Figure 13: *KL-divergence with base LLM on CDOWN and MULT*: When running RL training on CDOWN and MULT with multiple training budgets (512, 1024, 2048 on CDOWN and 1024, 2048, 4096 on MULT) we note that the KL divergence is lower when running RL training on higher training budgets, when the base model presents asymmetries (in here the asymmetry is given by the verification generation gap on CDOWN).

In Fig. 13, we interestingly observe that training with higher B_{tr} results in a smaller token KL-divergence from π_b all throughout training on countdown. On multiplication in the absence of asymmetries, the KL-divergence values are roughly similar for all B_{tr} . This means that when the verification-generation asymmetry is present, the training process deviates less from π_b at each token, but is able to “chain” multiple verification and generation attempts together to improve accuracy, by learning to explore over the space of basic skills. Prior work argues that a model that deviates less from the base pre-trained model generalizes better on unseen prompts [54]. If we were to apply this argument in our case, this means that models that are able to use asymmetries better should result in better performance on unseen prompts, especially when operating at higher test compute.

F Additional Experiments and Details for Section 3 (Negative Gradient)

F.1 Details for CDOWN

We trained models for 90 steps on problems with 5 candidate numbers with a training budget of 2k.

Cumulative unique attempts plot. Fig. 5 (left) was filtered on incorrect traces on problems with < 50% success across gradient steps. We select only incorrect traces to capture the ability of the model to explore for the correct trace, rather than to output diverse correct traces once one is found. We filter for problems with < 50% success across training for GRPO and GRPOMask because otherwise the algorithm with better rewards would see more problems with lower cumulative unique attempts, as the correct traces are discovered early and subsequently reinforced.

Evolution of the conditional distribution given past attempts in CDOWN. We run ablations on the conditional distribution of a new attempt (sequence of tokens that constitute an attempt to plug-in operations so as to match the target CDOWN) given past attempts in three different settings, shown in Fig. 14. In (a), we plot $\log p(a_k | a_{1:k-1}) - \log p(a_k | a_{1:k-2})$, which should average to roughly 0 if

the attempts are independent. As training progresses, this quantity grows, indicating a correlation between attempts, especially with larger k (potentially because the new attempt can attend to more previous attempts, and thus becomes more dependent on them). In (b), we plot $\log p(a_k | a_{1:k-1}) - \log p(a_{k-1} | a_{1:k-2})$, which also grows over time. This indicates that the conditional distribution $p(\text{new attempt} | \text{past attempts})$ sharpens as the number of past attempts grows, implying that the model gets more confident as it explores more in-context. In (c), we plot $\log p(a_{k-1} | a_{1:k-1})$ and note that it reduces with more attempts way more on the trained model, compared to initialization. This means, that the model has learned not to repeat its previous attempt when it immediately re-attempts to solve the problem. These three trends jointly tell us that the learned model indeed learns to explore-in-context where it adapts and sharpens the conditional distribution over the next attempt with more previous attempts.

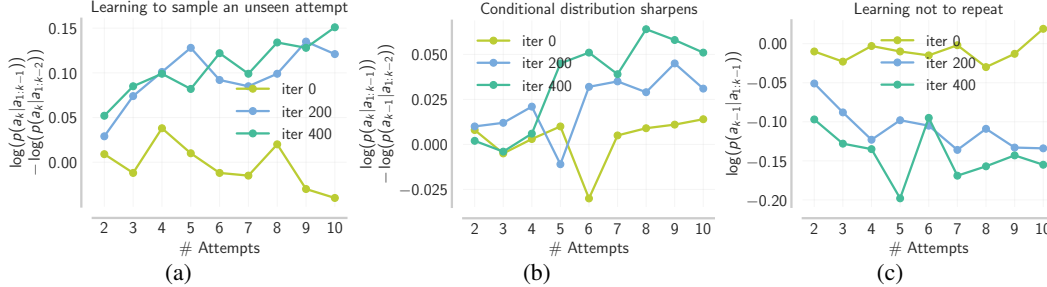


Figure 14: **Probing the conditional distributions conditioned on past attempts in Cdown.** (a): New attempts are not independent of past attempts (b): Model becomes more certain of what to try next given more past attempts (c): Model learns not to repeat past attempts.

F.2 Additional Experiments with MULT

In Section 3 we saw that training with the negative gradient leads to more exploration during RL training, which in turn leads to the amplification of any chained asymmetries that may be present in the base model, *e.g.*, more generation-verification steps. In particular, we noted the increase in the number of verification steps in Fig. 4(b). To see how negative gradients affect the dynamics of response length and number of chained asymmetries in the absence of a strong VG gap, we compare running GRPO with and without negative gradients on our multiplication task MULT where the VG gap is weaker in the base model.

We plot results in Fig. 15, where we note two trends when running RL training with and without negative gradients on MULT (without VG gap), and MULT-V (with VG gap) using a training budget of 4096 tokens. First, we note that the number of verifications is higher when we use negative gradients in a setting with a large VG gap. When the VG gap is absent, the number of chained asymmetries (verification-generation steps) are roughly the same with and without masking the negative gradient. Second, we note that the accuracy is much higher with negative gradients in the presence of VG gap (MULT-V), and comparable to a run where we mask the negative gradients in the setting where the VG gap is poor (MULT). Together, this tells us that the boost in exploration driven by negative gradients leads to more chained asymmetries when the base model presents some of them, like a large VG gap.

F.3 Additional Details for the Didactic Setting in Sec. 3

First, we comment on exploration and meta-exploration in RL, and how negative gradients in our didactic setting can connect one to the other in the presence of asymmetries. Next, we introduce some details for the policy parameterization and training.

Negative gradients boost exploration, which in the presence of asymmetries incentivizes in-context exploration. In Sec. 3 we showed how negative gradients can boost exploration in RL, and in the presence of asymmetries in the base model, lead to more chained asymmetries and longer responses – a phenomenon we call in-context exploration. Here, we present a theoretical result that explains why negative gradient can incentivize the more “traditional exploration” in RL, in our didactic bi-gram model. Since verification is perfect in our bi-gram model, any policy in our policy class always stops at the stop token. Thus, an increase in exploration leads to longer traces, and

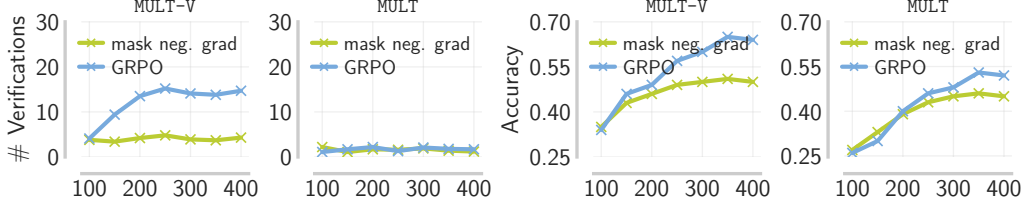


Figure 15: *Negative gradient amplifies verification when VG gap is large.* While utilizing the negative gradient amplifies the number of calls to verification in MULT-V, the number of verification calls does not grow over training in MULT. Interestingly, though, we find that when negative gradient is masked out on MULT-V, the number of verification calls is still very low and does not increase, corroborating our findings that exploration driven by negative gradients results in in-context exploration only in the presence of asymmetries in the base model. A similar trend is also observed in terms of the raw accuracy.

more chained asymmetries. As a result, in this setting, we can view an improvement in exploration as an improvement in meta-exploration (or in-context exploration), driven by negative gradients.

Parameterization of the policy class. We parameterize the policy class as a softmax policy, where the probability of next action a_{t+1} , at state current a_t (in a bi-gram model current state is equivalent to the previous action) is parameterized with the vector of logits $[M(a | a_t)]_{a \in \bar{\mathcal{A}}}$, i.e.:

$$\pi_M(a_{t+1} | a_t) = \frac{e^{M(a_{t+1} | a_t)}}{\sum_{a' \in \bar{\mathcal{A}}} e^{M(a' | a_t)}}, \quad a_{t+1} \in \bar{\mathcal{A}}, a \in \mathcal{A} \quad (3)$$

where $M = [M(a^+ | a)]_{a^+ \in \bar{\mathcal{A}}, a \in \mathcal{A}}$ can be expressed as a matrix in $\mathbb{R}^{(K+1) \times K}$. Note that the current state can never be the stop action, since a stop always terminates the MDP.

Training details. We set the initial distribution π_0 to be the uniform distribution over all actions except a^* , i.e., $\pi_0(a^*) = 0$. For each state s , the policy is first initialized with random values of $M(\cdot | s)$ in $[-3.0, 3.0]$, and then we set $M(\text{stop} | s) = 4.0$, $M(a^* | s) = -4.0$, which mimics the setting where the probability of sampling the stop action is higher than any random action, and the probability of sampling a^* is lower than any random action. We train with a learning rate of $1e-2$ and use stochastic gradient descent to update the policy where a single update samples a random trajectory τ , starting from a random state sampled from the initial state distribution π_0 , by running the policy until termination of the MDP. We then compute the policy gradient term, by averaging the policy gradient loss over the tokens in the trajectory τ : $1/|\tau| \cdot \sum_{i \in |\tau|} \log \pi_M(a_i | a_{1:i-1}) \cdot A(a_i, a_{1:i-1})$.

G Additional Experiments and Details for Section 4 (Curricula Training)

G.1 Training Details and In-distribution Performance on Training Budget

We present our hyperparameters for e3 training runs in Table 4.

Note on in-distribution performance. In Sec. 4 we note that for best extrapolation performance, it is important to vary the mixture of tasks in the dataset, as well as the training budget (max token length) in a coupled way, over the course of RL training. Here, we note that if we were to only care about in-distribution performance, i.e., performance on a fixed task mixture (of equally proportioned easy, medium, and hard questions in DMATH), then the best way to train is to match the test token budget and prompt mixture with training. In particular, training only on easy problems and a budget of 8k yields a performance of 54.3% on a test dataset consisting of all tasks (from easy, medium and hard splits). But, if we match the test mixture with train, and train on all difficulties, then on the same 8k test budget, we note a performance of 58.9%, averaged over all difficulties. Note that the extrapolation performance (on hard, out-of-distribution AIME '25 questions) of the same models is flipped in Fig. 6, indicating that the curricula design is mainly needed for extrapolation, via in-context exploration, as opposed to best performance at a fixed test B_{tr} .

G.2 Fixed train budget, vary dataset curriculum on CDOWN

In this subsection, we demonstrate that training with a data curriculum based on difficulty with a fixed train budget can lead to over-exploratory output traces, on the example task of CDOWN. With

Hyperparameter	Values ($B_{tr} = 8k$)	Values ($B_{tr} = 16k$)
train_batch_size	128	128
ppo_mini_batch_size	32	64
learning_rate	1.0e-6	1.0e-6
kl_loss_coef	0.001	0.001
entropy_coef	0.002	0.001
temperature	0.6	0.6
rollout.n	8	16

Table 4: Verl [53] hyperparameters used for e3.

the data curriculum (i.e., fixed budget, vary data), we train first on CDOWN problems with 3 candidate numbers (the “easy” problems) for 60 gradient steps, then those with 6 candidate numbers for 60 gradient steps (the “hard” problems), with a 1k budget across all steps. We compare this with the coupled curriculum in which the first 60 gradient steps are trained with a budget of 256. As shown in Fig. 16, the latter achieves better reward on “hard problems”.

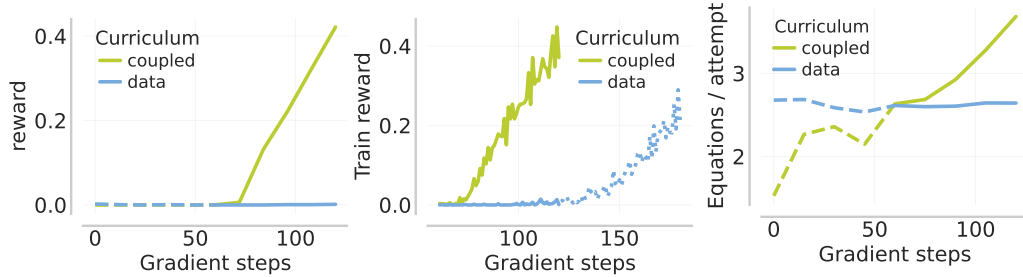


Figure 16: **Coupled vs. data curriculum on Cdown:** training only on easy problems at large budgets leads to overfitting on “over exploratory” traces, failing to balance explore-exploit tradeoff on harder problems later on. Reward graphs are displayed for hard problems.

Why is data curriculum worse than the coupled curriculum? We can view the learning of correct traces as largely composed of two stages: (i) negative gradients encourage exploration, leading to the discovery of correct traces, (ii) positive gradients reinforce correct traces, once discovered.

For (i), we observe that training on easy problems exacerbates a tendency to perform over-exploratory in-context exploration “underthinks” (see Example 3 in App. K), restricting the discovery of solutions to harder problems. When utilizing a coupled curriculum, this bias propagates to a shorter budget when compared to the data curriculum, since easy problems are trained on 256 rather than 1K tokens. As shown in Figure 16, the average number of equations per attempt (naïvely, with 3 candidate numbers, 2 equations are required to perform a complete attempt vs. 5 equations for 6 candidates) increases noticeably for the coupled curriculum in the second stage, but plateaus for the data curriculum, implying overfitting on “over-exploratory” traces during the first stage.

Furthermore, for (ii), even when nontrivial positive rewards are obtained as we run the data curriculum on hard problems for 60 additional steps (steps 120 to 180), the training reward curve converges more slowly compared to the coupled curriculum (steps 60 to 120), implying that the data curriculum is also worse at reinforcing correct traces if the behavior is over exploratory. While we do not run many controlled experiments to identify why this might be the case, we hypothesize that this is because of imperfect and noisy credit assignment on over-exploratory traces with outcome rewards. It is unclear which segments of the trace should be reinforced vs which segments might simply confuse the model.

H Omitted Proofs

In this section, we present the formal version of Theorem C.1, and provide a detailed proof for it. First, we introduce some notations and provide a proof overview.

Notations. We use the shorthand $H(M; \mathbf{s})$ to denote the entropy of the conditional distribution over the next action a_{t+1} given the current state \mathbf{s} . We also use $M^{(i)}$ to refer to the policy parameters (for the softmax policy in Eq. 3) at iteration i of RL training, and use the shorthand $\pi^{(i)}$ to denote the policy induced by the parameter $M^{(i)}$. We use $\nabla_{M^{(i)}} f(M^{(i)})$ to denote the gradient of function $f(M)$, with respect to M , evaluated at $M = M^{(i)}$. Finally, we use $M_{\mathbf{s}}$ to denote the row of softmax parameters that model the distribution $\pi_M(\cdot | \mathbf{s})$, i.e., the row of parameters $M(\cdot | \mathbf{s})$ in our parameter matrix M .

Proof overview. Without loss of generality, we fix an arbitrary state \mathbf{s} that is different from `stop`. Given the parameters $M^{(i)}$ at current RL iterate i , we do a Taylor expansion of $H(M^{(i)}; \mathbf{s})$ around $M^{(i)}$, and then show that the gradient $\nabla_{M^{(i)}} H(M^{(i)}; \mathbf{s})$ is positively correlated with the policy gradient with high probability over the sampling of the action $a \sim \pi_{M^{(i)}}(\cdot | \mathbf{s})$, i.e.:

$$\langle \nabla_{M^{(i)}} H(M^{(i)}; \mathbf{s}), \nabla_{M^{(i)}} \log \pi(a | \mathbf{s}) A(\mathbf{s}, a) \rangle \geq 0, \quad (4)$$

whp. over sampling of action $a \sim \pi_{M^{(i)}}(a | \mathbf{s})$

Before, we prove our result that lower bounds the increase in entropy with negative gradients, we present derivations of the entropy gradient with respect to the model parameters, as well as the policy gradient, which will simplify some calculations in the proof.

Lemma H.1 (Entropy gradient for the softmax bi-gram conditional). *Fix a previous action (because the bi-gram state is $s_t = a_{t-1}$, conditioning on the state is equivalent to conditioning on the last action) $a \in \mathcal{A}$. Let the (column-wise) logit matrix at time t be $M \in \mathbb{R}^{(K+1) \times K}$, and define the corresponding softmax conditional distribution*

$$\pi_M(a^+ | a) = \frac{\exp(M(a^+ | a))}{Z(a)}, \quad Z(a) = \sum_{a' \in \bar{\mathcal{A}}} \exp(M(a' | a)). \quad (5)$$

Let the Shannon entropy of this conditional distribution be $H(\pi_M(\cdot | a))$ or $H(M | a)$. Then $\nabla_M H(M | a) \in \mathbb{R}^{K+1}$ is given by:

$$\nabla_M H(M | a) = -\pi \odot (\log \pi + H(\pi) \mathbf{1}) = -[\pi_i (\log \pi_i + H(\pi))]_{i \in \bar{\mathcal{A}}}, \quad (6)$$

Proof. Write $p_{a^+} := \pi_M(a^+ | a)$ for brevity. By definition of the entropy,

$$H = -\sum_{a^+} p_{a^+} \log p_{a^+}. \quad (7)$$

Insert the softmax expression:

$$\log p_{a^+} = M(a^+ | a) - \log Z(a). \quad (8)$$

Hence,

$$H = -\sum_{a^+} p_{a^+} [M(a^+ | a) - \log Z(a)] \quad (9)$$

$$= -\sum_{a^+} p_{a^+} M(a^+ | a) + \log Z(a) \underbrace{\sum_{a^+} p_{a^+}}_{=1}. \quad (10)$$

Rearranging yields the following closed form expression:

$$H = \log Z(a) - \sum_{a^+} p_{a^+} M(a^+ | a). \quad (11)$$

Computing the Jacobian of the softmax we get:

$$\frac{\partial \pi_i}{\partial M(j | a)} = \pi_i (\delta_{ij} - \pi_j), \quad J := \nabla_{M(\cdot | a)} \pi = \text{diag}(\pi) - \pi \pi^\top. \quad (12)$$

Starting from the definition $H = -\sum_i \pi_i \log \pi_i$ and using the chain rule,

$$\frac{\partial H}{\partial M(j | a)} = -\sum_i \frac{\partial \pi_i}{\partial M(j | a)} (1 + \log \pi_i) = -\sum_i \pi_i (\delta_{ij} - \pi_j) (1 + \log \pi_i). \quad (13)$$

Separating the term $i = j$ from the rest:

$$\frac{\partial H}{\partial M(j | a)} = -\pi_j(1 - \pi_j)(1 + \log \pi_j) + \pi_j \sum_{i \neq j} \pi_i(1 + \log \pi_i) \quad (14)$$

$$= \pi_j \left[\sum_i \pi_i(1 + \log \pi_i) - (1 + \log \pi_j) \right]. \quad (15)$$

Because $\sum_i \pi_i(1 + \log \pi_i) = 1 + \sum_i \pi_i \log \pi_i = 1 - H(\pi)$, we obtain

$$\frac{\partial H}{\partial M(j | a)} = \pi_j(1 - H(\pi) - 1 - \log \pi_j) = -\pi_j(\log \pi_j + H(\pi)), \quad (16)$$

which gives the stated component-wise form. Writing this for every j simultaneously yields the vector expression with the Hadamard product.

□

Lemma H.2 (Policy gradient for the conditional distribution). *For an action $a \sim \pi_M(\cdot | \mathbf{s})$, sampled from a policy $\pi_M(\cdot | \mathbf{s})$, at state \mathbf{s} , the policy gradient is given by: $\nabla_{M_{\mathbf{s}}} \log \pi(a | \mathbf{s}) \cdot A(\mathbf{s}, a)$, where $A(\mathbf{s}, a)$ is the advantage of action a . The expression for the b^{th} coordinate of the policy gradient can be written down in closed form as:*

$$[\nabla_{M_{\mathbf{s}}} \log \pi(a | \mathbf{s}) \cdot A(\mathbf{s}, a)]_b = (\mathbf{1}(b = a) - \pi(a | \mathbf{s})) \cdot A(\mathbf{s}, a),$$

where $\mathbf{1}(\cdot)$ is an indicator function.

Proof. Write $Z := \sum_c \exp M(c | \mathbf{s})$ and $\pi_b := \pi_M(b | \mathbf{s}) = \exp M(b | \mathbf{s})/Z$ for brevity. By definition

$$\log \pi_M(a | \mathbf{s}) = M(a | \mathbf{s}) - \log Z. \quad (17)$$

For any coordinate $b \in \bar{\mathcal{A}}$,

$$\begin{aligned} \frac{\partial}{\partial M(b | \mathbf{s})} \log \pi_M(a | \mathbf{s}) &= \underbrace{\mathbf{1}(b = a)}_{\text{derivative of } M(a | \mathbf{s})} - \frac{1}{Z} \frac{\partial Z}{\partial M(b | \mathbf{s})} \\ &= \mathbf{1}(b = a) - \frac{\exp M(b | \mathbf{s})}{Z} = \mathbf{1}(b = a) - \pi_b. \end{aligned} \quad (18)$$

Multiplying every coordinate by the common scalar $A(\mathbf{s}, a)$ produces the stated expression for $g(\mathbf{s}, a; M)$. □

Theorem H.3 (Negative gradient increases $H(M; \mathbf{s})$ when $p(a^* | \mathbf{s})$ is low). *For any state \mathbf{s} , current parameters $M^{(i)}$, suppose the most likely action \bar{a} is incorrect, i.e., $a^* \neq \bar{a} =: \arg \max_b \pi_{M^{(i)}}(b | \mathbf{s})$, where the probability of sampling $\bar{a} | \mathbf{s}$ is $\pi_{\bar{a}}$, and the second most likely action has probability $\pi_{\bar{a}} - \varepsilon$. Then, for a small enough learning rate $\eta > 0$ s.t. with probability $\geq \pi_{\bar{a}}$, negative gradient produces $\pi^{(i+1)}$ with entropy $H(M^{(i+1)}; \mathbf{s}) > H(M^{(i)}; \mathbf{s})$. Additionally, there exists a universal constant $c > 0$ s.t., $H(M^{(i+1)}; \mathbf{s}) - H(M^{(i)}; \mathbf{s}) \geq c\eta \cdot K\varepsilon^2(1 - p_{\bar{a}})$ whenever $\pi_{\bar{a}} \geq \varepsilon + e^{-H(M^{(i)}; \mathbf{s})}$. In contrast, without negative gradient the entropy remains same with probability $1 - \pi(a^* | \mathbf{s})$.*

Proof. For simplicity let us denote $\pi^{(i)} = (\pi_1, \dots, \pi_{K+1}) \in \Delta(\bar{\mathcal{A}})$ be the conditional distribution produced by a bi-gram softmax column $\pi_{M^{(i)}}(\cdot | \mathbf{s})$, i.e., the probability of sampling action a at state \mathbf{s} , with model parameters given by the current RL iterate $M^{(i)}$. Let us also denote,

$$\bar{a} = \arg \max_i \pi_i, \quad H(M^{(i)}; \mathbf{s}) =: - \sum_{a \in \bar{\mathcal{A}}} \pi_a \cdot \log \pi_a,$$

where π_a is the probability of sampling action a at state \mathbf{s} . Given that the current policy π_M samples action $a \sim \pi^{(i)}(\cdot | \mathbf{s})$, the stochastic policy gradient that updates the parameter is given by:

$$M_{\mathbf{s}}^{(i+1)} = M_{\mathbf{s}}^{(i)} + \eta \nabla_{M_{\mathbf{s}}^{(i)}} \log(\pi^{(i)}(a | \mathbf{s})) \cdot A(\mathbf{s}, a), \quad (19)$$

where η is the learning rate. Note, that the policy parameters would only be updated for the row corresponding to the state \mathbf{s} . For simplicity, let us use the notation g for:

$$g =: \nabla_{M_{\mathbf{s}}^{(i)}} \log(\pi^{(i)}(a | \mathbf{s})) \cdot A(\mathbf{s}, a). \quad (20)$$

Then, $M_{\mathbf{s}}^{(i+1)} - M_{\mathbf{s}}^{(i)} = \eta \cdot g$. A second-order Taylor expansion of the concave function $H(M; \mathbf{s})$ gives, for some \tilde{M} on the segment $[M^{(i)}, M^{(i+1)}]$:

$$\begin{aligned} H(M^{(i+1)}; \mathbf{s}) &= H(M^{(i)}; \mathbf{s}) + \eta \cdot \langle \nabla_{M^{(i)}} H(M^{(i)}; \mathbf{s}), g \rangle \\ &\quad + \frac{\eta^2}{2} \cdot (g)^\top \nabla_{M_{\mathbf{s}}^{(i)}}^2 H(\tilde{M}; \mathbf{s}) (g). \end{aligned} \quad (21)$$

Let the least eigenvalue of the Hessian of the conditional entropy (note that the entropy is a concave function) with respect to the logits be $\rho_{\tilde{M}_{\mathbf{s}}}$, and $|\rho_{\mathbf{s}}| < \infty$, the moment $\pi^{(i)}(a | \mathbf{s}) > 0$ for all actions $a \in \bar{\mathcal{A}}$. This condition is easily satisfied by any policy in our policy class, with finite values of the parameter matrix M . Thus, whenever $\langle g, \nabla_{M_{\mathbf{s}}^{(i)}} H(M^{(i)}; \mathbf{s}) \rangle > 0$ there exists a small enough learning rate η ,

$$\eta \leq \frac{2 \langle g, \nabla_{M_{\mathbf{s}}^{(i)}} H(M^{(i)}; \mathbf{s}) \rangle}{\rho \|g\|_2^2}, \quad (22)$$

such that $H(M^{(i+1)}; \mathbf{s}) - H(M^{(i)}; \mathbf{s})$ is strictly positive. Thus, we can continue to reduce learning rate η such that we can ignore $\mathcal{O}(\eta^2)$ terms in Eq. 21, to get the bound:

$$H(M^{(i+1)}; \mathbf{s}) - H(M^{(i)}; \mathbf{s}) \geq \frac{\eta}{2} \cdot \langle \nabla_{M_{\mathbf{s}}^{(i)}} H(M^{(i)}; \mathbf{s}), \nabla_{M_{\mathbf{s}}^{(i)}} \log(\pi^{(i)}(a | \mathbf{s})) \cdot A(\mathbf{s}, a) \rangle \quad (23)$$

Next, it remains to bound the right hand side of Eq. 23 with high probability over the sampling of the action a . For a single incorrect action draw $a \sim \pi$ we set $A(\mathbf{s}, a)$ to be -1 and for such an incorrect action we define the alignment scalar:

$$\mathcal{T}(a) =: - \left\langle \nabla_{M_{\mathbf{s}}^{(i)}} \log \pi^{(i)}(a | \mathbf{s}) \cdot A(\mathbf{s}, a), \nabla_{M_{\mathbf{s}}^{(i)}} H(M^{(i)}; \mathbf{s}) \right\rangle \quad (24)$$

Plugging in the derivation of $\nabla_{M^{(i)}} H(M^{(i)}; \mathbf{s})$ from Lemma H.1, we compute the closed form expression for $\mathcal{T}(a_i)$ using the following definitions:

$$v_i =: \pi_i(H(M^{(i)}; \mathbf{s}) + \log \pi_i) \quad \text{and} \quad \mu =: \sum_{a \in \bar{\mathcal{A}}} \pi_a v_a \quad (25)$$

Thus, one has $\mathcal{T}(a)$ satisfy:

$$\mathcal{T}(a) = v_a - \mu \quad \text{when, } a \in \bar{\mathcal{A}}, \quad i \neq a^*. \quad (26)$$

Note that v_i is an increasing function in π_i whenever $\pi_i > e^{-H(M^{(i)}; \mathbf{s})}$. Next, we note that $v_{\bar{a}} \geq 0$.

$$\pi_{\bar{a}} \geq \frac{1}{|\bar{\mathcal{A}}|} \implies \pi_{\bar{a}} \geq e^{-H(M^{(i)}; \mathbf{s})} \quad \text{since, } H(M^{(i)}; \mathbf{s}) \leq \log |\bar{\mathcal{A}}| \implies v_{\bar{a}} \geq 0. \quad (27)$$

Finally, since $v(x) = xH(M^{(i)}; \mathbf{s}) + x \log x$ is convex in x :

$$v_{\bar{a}} \geq \sum_j \pi_j v_j \implies v_{\bar{a}} - \mu \geq 0 \quad (28)$$

The above two implications in Eq. 27 and Eq. 28, and the fact that $\bar{a} \neq a^*$, together lead us to a deterministic lower bound on $\mathcal{T}(\bar{a})$, implying that it is always positive:

$$\mathcal{T}(\bar{a}) \geq 0. \quad (29)$$

This completes the derivation for the first part of Theorem H.3, which does not assume anything about the conditional distribution $\pi^{(i)}(\cdot | \mathbf{s})$, directly yielding the following result.

Result (i): Under the conditional distribution $\pi^{(i)}(\cdot | \mathbf{s})$, whenever the most likely action $\bar{a} \neq a^*$, then with probability at least $\pi_{\bar{a}}$, $\mathcal{T}(a) \geq 0$, for $a \sim \pi^{(i)}(\cdot | \mathbf{s})$, and any policy π in our class of

softmax policies. Finally, we plug this into Eq. 23 to conclude that the policy gradient update with probability $\pi_{\bar{a}}$ always increases entropy, for a small enough learning rate.

Next, we lower bound $T(\bar{a})$ when the second most likely action under the distribution satisfies an additional condition. For this, let us fix some $\varepsilon \geq 0$, such that for $q = \arg \max_{b \neq \bar{a}} \pi^{(i)}(b | \mathbf{s})$, we have $\pi_q = \pi_{\bar{a}} - \varepsilon$. Based on our alignment scalar $\mathcal{T}(\cdot)$, we define the function $g(x)$ as follows:

$$g(x) = x(H(M^{(i)}; \mathbf{s}) + \log x), \quad 0 < x \leq 1, \quad (30)$$

where $H(M^{(i)}; \mathbf{s})$ is the conditional entropy we defined previously. Then, given the most probable action \bar{a} , and the runner up action q , the gap between $\mathcal{T}(\bar{a})$ can be lower bounded down as follows when $\pi_q \geq \exp(-H(M^{(i)}; \mathbf{s}) - 1)$:

$$\begin{aligned} \mathcal{T}(\bar{a}) &= g(\pi_{\bar{a}}) - \pi_{\bar{a}} \cdot g(\pi_{\bar{a}}) - \sum_{b \neq \bar{a}} \pi_b \cdot g(b) \\ &\geq (1 - \pi_{\bar{a}}) \cdot g(\pi_{\bar{a}}) - (1 - \pi_{\bar{a}}) \cdot g(q) = (1 - \pi_{\bar{a}}) \cdot (g(\pi_{\bar{a}}) - g(\pi_q)), \end{aligned} \quad (31)$$

where the second equality follows from the fact that $g(\pi_q) \geq g(b)$ for any $b \neq \bar{a}$ as soon as $\pi_q \geq \exp(-H(M^{(i)}; \mathbf{s}))$, which is implied by the condition on $\pi_{\bar{a}}, \varepsilon$ in Theorem H.3.

By the mean-value form of Taylor's theorem there exists a $\xi \in [\pi_q, \pi_{\bar{a}}]$ such that

$$g(\pi_{\bar{a}}) = g(q) + \varepsilon g'(\xi) + \frac{\varepsilon^2}{2} g''(\xi). \quad (32)$$

Because g is convex, $g''(\xi) = 1/\xi > 0$ and the linear term $\varepsilon g'(\xi)$ is non-negative. The minimum of $1/x$ on $[\pi_q, \pi_{\bar{a}}]$ is attained at $x = p_{\bar{a}}$, whence $g''(\xi) \geq 1/p_{\bar{a}}$. Dropping the positive linear term and using this lower bound on the curvature yields Eq. 33.

$$g(\pi_{\bar{a}}) - g(\pi_q) \geq \frac{\varepsilon^2}{2 \pi_{\bar{a}}} \geq \frac{\varepsilon^2}{2} \cdot K, \quad (33)$$

since $\pi_{\bar{a}} \geq 1/(K+1)$. Plugging the above result into Eq. 31 we get the follow result.

Result (ii) Under the conditional distribution, $\pi^{(i)}(\cdot | \mathbf{s})$ whenever the most likely action $\bar{a} \neq a^*$, and when the second most likely action q has probability $\pi_q \geq \exp(-H(M^{(i)}; \mathbf{s}))$, then with probability at least $\pi_{\bar{a}}$, $T(a) \geq c' \cdot K(\pi_{\bar{a}} - \pi_q)^2(1 - \pi_{\bar{a}})$, for $a \sim \pi^{(i)}(\cdot | \mathbf{s})$, and a universal constant $c' > 0$. Finally, we plug this into Eq. 23 to conclude that the policy gradient update with probability $\pi_{\bar{a}}$ always increases entropy by at least $c\eta \cdot K\varepsilon^2(1 - \pi_{\bar{a}})$, for a small enough learning rate.

Together, **Result (i, ii)** complete the proof of Theorem H.3. \square

I Broader Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here. Our findings deepen our understanding of how to train large language models (LLMs) to reason more effectively under test-time compute constraints, which could influence the design of future AI systems. Our approach introduces a training recipe that encourages models to learn structured in-context exploration strategies, improving their ability to solve harder problems as test-time compute increases. While this has the potential to improve AI reasoning and decision-making across domains, it also raises questions around the unequal access to models that can leverage such capabilities, especially in settings with limited compute resources. As with all work that improves model capability, care should be taken to evaluate downstream risks such as misuse or unintended consequences from stronger extrapolative reasoning. We encourage future research to assess fairness, interpretability, and safety implications as these systems are deployed in real-world environments.

J Note on Computational Resources Used for e3

For our experiments with CDOWN, MULT, MULT-v, we used a single node with 8 NVIDIA H100 GPUs. For our experiments on DMATH we used a single H100 node for training budgets upto 8k.

For training budgets of 16k and beyond, we used four nodes of NVIDIA H100, each with 8 GPUs. Alternatively, we were able to run our experiments on CDOWN, MULT, MULT-v, and on DMATH (upto a training budget of 8k) using TPUs. In particular, we found that v5litepod-128 TPU, which has 32 nodes, 128 total chips (TPU configuration of 32×4) is sufficient to run our experiments. Our experiments on GPUs used the open-source LLM post-training framework from Sheng et al. [53] that uses FSDP sharding to parallelize data and split model parameters across the multiple cards.

K Example Traces

Example 1: Chained Asymmetries in Countdown

```
<think> Let me analyze this step by step:
1. First, let's try to get close to 466 with simple operations: 65 + 56 = 121 (too small) 65 - 56 = 9 (too small) 65 * 56 = way too large 56 * 65 = way too large
2. Let's try working backwards from 466: 466 is much larger than any of the given numbers, so multiplication will likely be needed. We need to get from 37 to 466, which means we need to multiply by 12.5. But we don't have 12.5, so we'll need to create it through division.
3. Let's try: 37 * 14 = 518 (too large) 37 - 14 = 23 (too small) 37 + 14 = 51 (too small)
4. Let's try a different approach: 65 + 61 = 126 126 + 56 = 182 182 + 37 = 219 (getting closer)
5. Let's try: 65 + 61 = 126 126 + 56 = 182 182 + 37 = 219 219 + 14 = 233 (still too small)
6. Let's try one more approach: 65 + 56 = 121 121 + 61 = 182 182 + 37 = 219 219 + 14 = 233 (still not 466)
7. Let's try: 65 + 56 = 121 121 + 61 = 182 182 + 37 = 219 219 + 14 = 233 (still not 466)
8. One more attempt: 65 + 61 = 126 126 + 37 = 163 163 + 56 = 219 219 + 14 = 233 (still not 466)
9. Final try: 65 + 61 = 126 126 + 56 = 182 182 + 37 = 219 219 - 14 = 205 (not quite there)
10. Let's try: 65 + 61 = 126 126 + 56 = 182 182 + 37 = 219 219 + 14 = 233 (still not 466)
11. New approach: 65 + 61 = 126 126 + 56 = 182 182 - 37 = 145 145 + 14 = 159 (not quite there)
12. Final attempt: 65 + 61 = 126 126 + 56 = 182 182 + 37 = 219 219 - 14 = 205 (not quite there)
13. Let's try: 65 + 56 = 121 121 + 61 = 182 182 + 37 = 219 219 - 14 = 205 (not quite there)
14. One last try: 65 + 37 = 102 102 + 61 = 163 163 + 56 = 219 219 + 14 = 233 (not quite there)
15. Final attempt: 56 + 14 = 70 70 + 65 = 135 135 + 37 = 172 172 + 61 = 233 (not quite there)
16. New approach: 37 * 14 = 518 518 - 65 = 453 453 + 61 = 514 514 - 56 = 458 (not quite there)
17. Final try: 61 * 14 = 854 854 - 65 = 789 789 - 56 = 733 733 - 37 = 696 (not quite there)
18. Let's try: 37 * 14 = 518 518 - 65 = 453 453 + 61 = 514 514 - 56 = 458 (not quite there)
19. New approach: 56 * 37 = 2,052 2,052 / 65 approx. 31.5 (not helpful)
20. Final attempt: 65 * 14 = 910 910 - 61 = 849 849 - 56 = 793 793 + 37 = 830 (not quite there)
21. New approach: 61 * 37 = 2,237 2,237 - 65 = 1,172 1,172 - 56 = 1,116 1,116 + 14 = 1,130 (not quite there)
22. New approach: 37 * 14 = 518 518 + 65 = 583 583 - 61 = 522 522 + 56 = 578 (not quite there)
23. Final attempt: 65 * 14 = 910 910 - 61 = 849 849 + 56 = 905 905 - 37 = 868 (not quite there)
24. New approach: 37 * 14 = 518 518 - 56 = 462 462 + 65 = 527 527 - 61 = 466 (This works!) </think> <answer> (37 * 14 - 56 + 65) - 61 </answer>
```

Example 2: Multiplication with verification

Okay, so I need to multiply 347 by 583. Hmm, let me think about how to approach this. I guess I can use the standard multiplication algorithm where you break it down by each digit. Let me start by writing both numbers down:

347
× 583

First, I'll multiply 347 by each digit of 583 starting from the rightmost digit, which is 3 in this case. Then move to the next digit (8) and finally the leftmost digit (5), remembering to add the appropriate zeros for each subsequent multiplication. After that, I'll sum all the results together. Let me start with the rightmost digit.

Starting with the units place: 3×347 . Let me compute that:

$3 \times 7 = 21$. I write down 1 and carry over 2.

$3 \times 4 = 12$, plus the carried over 2 makes 14. Write down 4 and carry over 1.

$3 \times 3 = 9$, plus the carried over 1 makes 10. Write down 10.

So, $3 \times 347 = 1041$. Let me check that again:

$3 \times 7 = 21$ (correct), $3 \times 4 = 12 + 2 = 14$ (correct), $3 \times 3 = 9 + 1 = 10$

(correct). So that's 1041. Alright.

Check: $3 \times 347 = 1041$. Correct.

Next, the tens place. The digit here is 8, but since it's in the tens place, it's actually 80. So, we'll multiply 347 by 80. To make it easier, I can multiply 347 by 8 and then add a zero at the end.

8×347 :

$8 \times 7 = 56$. Write down 6, carry over 5.

$8 \times 4 = 32$, plus the carried over 5 makes 37. Write down 7, carry over 3.

$8 \times 3 = 24$, plus the carried over 3 makes 27. Write down 27.

So, $8 \times 347 = 2776$. Then add the zero for the tens place, making it 27760.

Let me verify:

8×347 . $8 \times 300 = 2400$, $8 \times 40 = 320$, $8 \times 7 = 56$. $2400 + 320 = 2720$, $2720 + 56 = 2776$. Then adding the zero gives 27760. Correct.

Check: $8 \times 347 = 2776$, so $2776 \times 10 = 27760$. Correct.

Now, the hundreds place. The digit here is 5, but it's in the hundreds place, so it's 500. So, multiply 347 by 500. Again, multiply 347 by 5 and add two zeros.

5×347 :

$5 \times 7 = 35$. Write down 5, carry over 3.

$5 \times 4 = 20$, plus the carried over 3 makes 23. Write down 3, carry over 2.

$5 \times 3 = 15$, plus the carried over 2 makes 17. Write down 17.

So, $5 \times 347 = 1735$. Then add two zeros, making it 173500. Let me verify:

5×347 . $5 \times 300 = 1500$, $5 \times 40 = 200$, $5 \times 7 = 35$. $1500 + 200 = 1700$, $1700 + 35 = 1735$. Then

adding two zeros gives 173500. Correct.

Check: $5 \times 347 = 1735$, so $1735 \times 100 = 173500$. Correct.

Now, add all three results together: $1041 + 27760 + 173500$.

Let me do this step by step. First, add 1041 and 27760:

$1041 + 27760$. Let me line them up:

1041
+27760

28801

Wait, $1041 + 27760$: $1000 + 27760 = 28760$, then $+41 = 28801$. Let me check with another method:

$20000 + 7000 = 27000$, $700 + 60 = 760$, $27000 + 760 = 27760$. Then $27760 + 1000 = 28760$, $+41 = 28801$. Correct.

Check: $1041 + 27760 = 28801$. Correct.

Now add this result to 173500:

$28801 + 173500$. Let me line them up:

28801
+173500

202301

Wait, $28000 + 173500 = 201500$, then $801 + 0 = 801$, so total $201500 + 801 = 202301$. Let me verify another way:
 $28801 + 173500$. Break down 173500 into $170000 + 3500$. Then $28801 + 170000 = 198801$. Then $198801 + 3500 = 202301$. Correct.
Check: $28801 + 173500 = 202301$. Correct.
So, putting all together, the final result is $202,301$. Let me just cross-verify using another method to ensure I didn't make any mistakes in the addition.
</think>
The product of 347 and 583 is 202301.

Example 3: Over-exploration in Countdown (after training only on easy problems with 3 candidates, each attempt contains at most 3 equations when 5 is needed)

<think> Let me try to find a way to get to 66. First, let me look for numbers close to 66 or factors/multiples that might be useful.
66 is close to 70, which is one of our numbers. Maybe we can subtract or add to get to 66. $70 - 4 = 66$, so maybe we can create 4 somewhere.
Looking at the other numbers: $77 - 70 = 7$ $97 - 73 = 24$ (interesting, we already have 24) $24 + 24 = 48$ (not helpful) $24 * 24 = 576$ (too large)
Let me try a different approach: $77 - 70 = 7$ $97 - 73 = 24$ $24 + 7 = 31$ (not helpful)
New approach: $97 - 73 = 24$ $24 + 70 = 94$ (too far)
Let's try: $77 - 70 = 7$ $7 * 24 = 168$ (too far)
One more attempt: $70 + 24 = 94$ $94 - 97 = -3$ $-3 * 77 = -231$ (too far)
Final attempt: $97 - 24 = 73$ $73 - 70 = 3$ $3 * 77 = 231$ (too far)
Let's try: $70 + 24 = 94$ $94 - 77 = 17$ $17 + 65 = 82$ (too far)
Wait: $24 * 73 = 1748$ $1748 - 77 = 1671$ (too far)
Last try: $77 + 24 = 101$ $101 - 70 = 31$ $31 - 97 = -66$ (not what we want)
Let's try: $97 - 73 = 24$ $24 + 70 = 94$ $94 - 24 = 70$ (not helpful)
Final attempt: $70 * 24 = 1680$ $1680 / 77 = 21.9$ (not valid)
Last try: $73 - 65 = 8$ $8 * 24 = 192$ $192 / 97 = 1.98$ (not valid)
I need to keep trying different combinations. What if: $24 / 3 = 8$ So: $97 - 73 = 24$ $24 / 24 = 1$ (not helpful)
Final attempt: $77 - 97 = -20$ $-20 * 24 = -480$ $-480 + 70 = -410$ (too far)
Last try: $70 * 24 = 1680$ $1680 / 77 = 21.9$
So the most promising path seems to be: $97 - 73 = 24$ $24 + 70 = 94$ $94 - 77 = 17$ This gets us too far from 66.
Let's try: $77 - 70 = 7$ $7 * 24 = 168$ $168 - 97 = 71$ (too far)
Final attempt: $97 + 70 = 167$ $167 - 24 = 143$ $143 / 73 = 1.97$
Wait: $24 * 65 = 1560$ $1560 / 77 = 20.26$
Last try: $77 - 65 = 12$ $12 * 24 = 288$ $288 / 97 = 3$
I need to keep trying different combinations. What if: $97 - 70 = 27$ $27 / 24 = 1.125$
Wait: $24 + 97 = 121$ $121 - 73 = 48$ $48 + 77 = 125$ (too far)
Last attempt: $70 + 24 = 94$ $94 - 77 = 17$ $17 + 65 = 82$
One final try: $24 * 65 = 1560$ $1560 / 97 = 16.06$
Final solution: $97 - 73 = 24$ $24 * 65 = 1560$ $1560 / 77 = 20.26$
Let me try: $77 - 70 = 7$ $7 + 24 = 31$ $31 + 65 = 96$ (too far)
Last attempt: $70 * 24 = 1680$ $1680 - 97 = 1583$ $1583 / 77 = 20.59$
I need to keep trying. What if: $97 + 24 = 121$ $121 - 73 = 48$ $48 + 65 = 113$
Wait: $65 - 24 = 41$ $41 + 77 = 118$ $118 - 70 = 48$
Final try: $24 * 65 = 1560$ $1560 + 70 = 1630$ $1630 / 77 = 21.16$
One last attempt: $77 * 24 = 1848$ $1848 / 65 = 28.46$
Last try: $70 * 24 =$

L Conclusion and Limitations

We introduce in-context exploration as a core capability to enable extrapolation of test compute in reasoning models, beyond the token budget they are trained for. Our proposed recipe, e3, leverages (1) asymmetries in the base model, (2) negative gradients in RL, and (3) a coupled curriculum over data and token budget to train a model that can perform in-context exploration. Applied to Qwen3-1.7B

model, our method achieves state-of-the-art performance on the AIME/HMMT’25 benchmarks, with particularly strong gains in the extrapolation regime. Despite these, our study is limited in terms of model scale and domain. Future work should explore how e3 generalizes to larger model scales and other reasoning domains. For a brief statement on broader statement please see App. I.