RL FOR REASONING BY ADAPTIVELY REVEALING RATIONALES

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

016

017

018

019

021

025

026

027

028

031

033

034

037

038

040

041

042

043

044

046 047

048

051

052

ABSTRACT

Learning in the combinatorially large output space of sequence generation problems is challenging as providing expert demonstrations scales poorly with sequence length, and RL struggles with sparse rewards. Between dense demonstrations in supervised training and no demonstrations in reinforcement learning lies an underexplored regime: partial supervision. We ask whether some classes of sequence learning problems become efficiently learnable by exploiting this gap. We address this by introducing adaptive backtracking (AdaBack), a per-sample curriculum learning algorithm that reveals a partial prefix of the target output. The supervision length is adjusted dynamically for each sample based on the model's past reward signal, allowing it to incrementally learn to complete reasoning chains by conditioning on correct partial solutions. We investigate this intermediate regime between SFT and RL and argue that per-sample curriculum learning is more than a trade-off between efficiency and generality—it can succeed in tasks with long sequences of latent dependencies where SFT and RL both fail to generalize. Using a synthetic task with latent parity constraints, we show that AdaBack reliably solves problems that are otherwise intractable. On two mathematical reasoning benchmarks, MATH and GSM8k, we find that AdaBack enables models to solve problems that RL alone cannot, acquiring new reasoning capabilities through incremental exposure to partial solutions.

1 Introduction

The reasoning capabilities of Transformers (Vaswani et al., 2017) have been extensively studied across domains such as mathematics (Saxton et al., 2019; Cobbe et al., 2021; Hendrycks et al., 2021; Lewkowycz et al., 2022), algorithmic reasoning (Veličković et al., 2022), and code generation (Chen et al., 2021). These studies indicate that the reasoning performance of such models is significantly enhanced by training and inference mechanisms involving explicit reasoning traces or rationales, commonly known as scratchpad (Nye et al., 2021) or chain-of-thought (Wei et al., 2023).

However, acquiring large-scale, high-quality reasoning traces for specialized domains, such as mathematics, poses substantial challenges. Reinforcement learning (RL)-inspired methods have emerged as a promising solution to this challenge. By utilizing (potentially noisy) reward functions or verifiers (e.g., checking final answers to math problems), language models can generate novel reasoning traces, effectively using the model as an RL policy. For instance, the REINFORCE algorithm (Williams, 1992) has been employed by methods like STaR (Zelikman et al., 2022), generating multiple solutions per problem and selectively fine-tuning on the correct ones. More advanced methods utilize algorithms like Proximal Policy Optimization (PPO) (Schulman et al., 2017) and, more recently, Group Relative Policy Optimization (GRPO) (Shao et al., 2024).

Despite recent progress, RL-based approaches for structured reasoning tasks continue to face significant hurdles. The key challenge lies in exploration: as reasoning chains grow longer, the space of valid output sequences increases exponentially, while reward signals remain sparse and often binary. Consequently, the probability of sampling a correct solution through random exploration diminishes exponentially with sequence length. As a result, standard RL tends to reinforce reasoning paths that are already assigned non-negligible probability by the pretrained model. Empirical evaluations by Havrilla et al. (2024) and Yue et al. (2025) support this observation, showing that RL fine-tuning primarily amplifies existing behaviors without substantially exploring novel solution trajectories.

These limitations motivate our investigation of a third regime: the space **between supervised fine-tuning (SFT) and RL**. We study whether breaking down long reasoning chains into partial rationales—and adaptively conditioning the model on only a prefix of the target solution during training—can make exploration tractable, and ultimately expand the class of problems that sequence models can learn to solve. This leads to our central research question:

Can reinforcement learning, when guided by adaptive partial supervision, teach models genuinely new reasoning capabilities? Specifically, for sequences learning tasks of variable lengths, can it enable discovery of solutions that were previously exponentially (in sequence length) unlikely under the model's initial distribution?

To illustrate this clearly, consider a reasoning task consisting of n sequential steps, each of which must be performed correctly for the task to succeed (e.g., proving a math theorem). For simplicity, assume the model performs each step correctly with a constant probability p. Naively attempting the entire task yields a success probability of p^n , meaning that positive reinforcement signals would be exponentially infrequent—on average, every p^{-n} iterations—making standard RL impractical. To mitigate this, we propose adaptive guidance: initially, we reveal all but the final step of the solution from the training dataset, using RL solely for this final step. Consequently, the model receives feedback with probability p every iteration. As performance improves, we progressively conceal more steps of the solution, training the model to complete increasingly larger segments, maintaining frequent positive feedback. Eventually, the entire reasoning chain is learned step-by-step, effectively transforming a complex search space with success probability p^n into n simpler sub-searches, each with success probability $\Theta(p)$. While this example includes simplifying assumptions, we empirically demonstrate this phenomenon's practicality on a parity task described in Section 2.2.

The scenario described above is simplified and idealized. In real-world problems, reasoning steps may not be clearly distinguishable, and datasets typically contain tasks with varying difficulties and solution lengths, making uniform partial reveals inefficient. To address these practical challenges, we propose a per-sample adaptive algorithm that dynamically adjusts the revealed portion of each solution based on its inherent difficulty. Specifically, we leverage the GRPO framework (Shao et al., 2024), which naturally estimates task difficulty by generating multiple rollouts per question and averaging the rewards. We term this method *adaptive backtracking* (*AdaBack*), detailed further in Section 2.

Backtracking strategies have precedents in RL—for example, the method of Salimans and Chen (2018). In language modeling, R3 (Xi et al., 2024) applies a curriculum that progressively trains models to complete reasoning chains from earlier positions. Building on these ideas, we ask: how can such slicing be adapted per sample, guided by model performance at each iteration, without relying on a heuristically designed global curriculum?

Contributions Our contributions are as follows:

- We propose AdaBack, a per-sample adaptive curriculum learning algorithm that dynamically
 adjusts the amount of supervision during RL training based on reward feedback. AdaBack enables
 each sample to progress at its own rate, eliminating the need for manual curriculum staging or
 handcrafted schedules, transitioning from full supervision to full exploration in a data-driven way.
- We show that AdaBack bridges the gap between SFT and RL, enabling learning in regimes where both fail. On a synthetic parity task with sparse rewards, we demonstrate a **separation result**: AdaBack reliably solves the task, while SFT, RL, and their combination all fail.
- On standard mathematical reasoning benchmarks (MATH, GSM8k), we show that **AdaBack improves performance over standard RL and SFT+RL pipelines**. We also show that AdaBack applied to base models often matches the performance of SFT-initialized counterparts.
- We introduce further two variants of GSM8k: **Base-7**, which uses an unfamiliar numerical format never seen by the model during pretraining, and **Tensor-2**, which concatenates problems to increase reasoning depth. AdaBack achieves strong performance on both, demonstrating **robust generalization to symbolic shifts and longer-horizon reasoning**.
- Finally, we identify a limitation: for instruct-tuned models or models where pretraining has already exposed the model to most problem types, AdaBack together with standard RL training provide no benefits, underscoring their limitation in aiding exploration where uncertainty is low.




Figure 1: AdaBack Update Rule. At epoch t, we sample a supervision ratio $\rho_t^{(i)} \sim U(\rho_{\min}^{(i)}, \rho_{\max}^{(i)})$ and condition the model's generation on the question and the corresponding partial answer (shown in black text, with unrevealed content grayed out). If the average reward is below a threshold τ , we increase supervision in the next epoch by sampling from the red interval $\rho_{t+1}^{(i)} \sim U(\rho_t^{(i)}, \rho_{\max}^{(i)})$. Otherwise, if the reward exceeds τ , we reduce supervision and sample from $U(0, \rho_t^{(i)})$, the green interval, to make the task harder.

2 ADAPTIVE BACKTRACKING

The core idea is to expose a prefix of the target sequence during training and gradually reduce this supervision based on model performance. Unlike fixed-step or handcrafted curricula, this method allows each sample to progress at its own pace, naturally balancing difficulty and learning progress.

For RL algorithms like GRPO (Shao et al., 2024) or RLOO (Ahmadian et al., 2024) that perform multiple rollouts for each sample, we dynamically adjust the supervision level so that the average sample reward across the rollouts stays close to a desired amount (for example, 50% of the answers are correct for each math question). For other RL algorithms, the criteria to change the supervision level can depend on other estimation of the state's value.

Problem Setup We denote random variables using capital letters. Let an input sample be represented as a sequence of tokens $X=(X_1,X_2,\ldots,X_\ell)$, corresponding to a natural language prompt, such as a question or problem description. The desired output is a chain-of-thought (CoT) style response $Y=(Y_1,Y_2,\ldots,Y_m)$, where each Y_t represents a token or step in the model's reasoning process. We assume we have a dataset of such (X,Y) pairs.

Our goal is to train a model to generate a correct reasoning sequence Y conditioned on input X. Importantly, the correct CoT is not necessarily unique: multiple reasoning paths may lead to valid answers. However, we assume that correctness is verifiable via a reward model $r = R(X_{1:\ell}, \hat{Y}_{1:m})$ where r=1 if the generated output \hat{Y} is accepted as correct. We denote by $r_{\text{format}} \in [0,1)$ the reward for a parsable generation when the generation has the correct format, and assume that r=0 otherwise. This setup reflects many real-world reasoning tasks. For instance, in mathematical problem solving, final answers can often be verified, and different CoT traces may yield the same final result.

2.1 Method

At each training step, given an input $X^{(i)} = (X_1^{(i)}, \dots, X_{\ell_i}^{(i)})$ and its corresponding target output $Y^{(i)} = (Y_1^{(i)}, \dots, Y_{m_i}^{(i)})$, we reveal the first k tokens of Y, $Y_{1:k}$, where $k = \lfloor \rho^{(i)} \cdot m_i \rfloor$ and $\rho^{(i)} \in [0,1]$ is a sample-specific supervision portion. The model is trained to continue generation conditioned on the input and the revealed prefix $\hat{Y}_{k+1:m_i'}^{(i)} \sim P_{\theta}(\cdot \mid X^{(i)}, Y_{1:k}^{(i)})$, where $m_i' - k$ is the length of the generated continuation.

Adaptive Update Rule For each training sample i, we maintain an interval $[\rho_{\min}^{(i)}, \rho_{\max}^{(i)}]$ initialized as [0,1], from which we uniformly sample the supervision portion $\rho_t^{(i)}$ for sample i at epoch t. After generating a set of predictions and receiving an average reward $r_t^{(i)}$ (or other estimations of the state's value if not using GRPO), we update this interval based on a fixed reward threshold τ :

$$\begin{split} & \text{If } r_t^{(i)} < \tau: \quad \rho_{\min}^{(i)} \leftarrow \rho_t^{(i)} \\ & \text{If } r_t^{(i)} \geq \tau: \quad \rho_{\max}^{(i)} \leftarrow \rho_t^{(i)}, \quad \rho_{\min}^{(i)} \leftarrow 0.0 \end{split}$$

The next level of supervision is then uniformly sampled from this updated range:

$$\rho_{t+1}^{(i)} \sim U(\rho_{\min}^{(i)}, \rho_{\max}^{(i)}).$$

This process is visualized in Figure 1.

Intuitively, the model receives less supervision when it performs well and more when it struggles, enabling an automatic per-sample curriculum that adapts to the model's learning progress. This procedure naturally drives the model toward completing longer and more challenging portions of the target sequence, only when it shows competence on easier prefixes. The threshold τ governs how strictly we evaluate model success: higher values require stronger reward signals (e.g., higher average rewards) for the supervision ratio to decrease. In general terms, this process performs a form of stochastic binary search over the supervision ratio ρ , using reward feedback as a success signal. The goal is to minimize the amount of revealed rationale while ensuring the model continues to receive useful reward signals. As the model improves, the supervision ratio naturally decreases. For samples with no reward history, we initialize $\rho^{(i)}$ from global moving averages $\bar{\rho}_{\min}$ and $\bar{\rho}_{\max}$, which are updated over time using exponential moving averages. Additionally, to close train-test distribution mismatch, with a small probability we randomly set the portion to zero. These training details are further discussed in Appendix C.

This adaptive scheme enables each training sample to **follow its own trajectory from full supervision to full generation,** providing a flexible and data-efficient approach to curriculum learning in structured sequence tasks. This per-sample scheduling ensures that each training point advances only when ready, allowing the model to incrementally acquire reasoning skills without overfitting to fixed patterns.

2.2 Chain-of-Parities: A Synthetic Environment for Studying Reasoning

To better understand learning dynamics in isolation from the complexities of natural language and pretraining, we introduce a synthetic sequence modeling task called the **chain-of-parities**. This task could be viewed as a **contextual blind cliff walk**, inspired by Schaul et al. (2015), adapted to reflect the challenges of chain-of-thought reasoning.

We aim to address a fundamental question: Are there sequence learning tasks that cannot be learned by SFT, RL, or their naive combination but that can be solved by Adaptive Backtracking? We construct such a task and demonstrate that the answer is affirmative.

Given a binary input sequence $X \in \{0,1\}^L$, the goal is to generate an output sequence $Y_1, Z_1, \dots, Y_L, Z_L$ of length 2L, where

- $Y_i \in \{0,1\}$ is unconstrained (both values are acceptable),
- Z_i is the parity of X_i , Y_i , and Z_{i-1} , i.e, $Z_i = Z_{i-1} \oplus Y_i \oplus X_i$, where \oplus denotes the XOR function and $Z_0 = 0$.

This design enforces a latent, step-wise structure over the output. The Y_i values act as a "scratch-pad"—arbitrary values that influence Z_i through an accumulating parity computation. In essence, each Z_i encodes the parity of the prefix $(X_1,\ldots,X_i,Y_1,\ldots,Y_i)$, making the correctness of Z_i dependent on the accuracy of Z_{i-1} and prior chain-of-thought (CoT) steps. This recursive dependency mirrors real-world CoT tasks, where early mistakes cascade through the solution. Although there are 2^L valid outputs per input due to unconstrained Y_i , generating one randomly has a probability of just 2^{-L} —analogous to sparse-reward regimes common in reasoning tasks.

We consider learning this task from a dataset of n uniformly sampled sequences (X_i, Y_i) . This setup highlights the inherent limitations of SFT and RL:

- **RL fails**: Rewards are sparse, and discovering a single valid output via random exploration becomes exponentially unlikely as the problem length grows.
- SFT fails: For small training sets, supervised training alone cannot learn this task. Specifically, the task involves learning n-1 parity functions of degree three, for which the theoretical sample complexity is well-studied (Abbe et al., 2023a; Kou et al., 2024). For example, in the statistical query (SQ) model (Kearns, 1998), weakly learning a degree-k parity from L input bits requires at

least $\Omega(L^{k-1})$ samples. Consequently, regular fixed-size neural networks trained via stochastic gradient descent (SGD) cannot even weakly learn a degree-three parity from only $n = \tilde{O}(L)^1$ samples (Abbe et al., 2023a).

• SFT + RL fails: With a limited number of training samples, such as $n = \tilde{O}(L)$ discussed above, SFT does not provide weak learning. Therefore, the exploration of the model remains random, and rewards remain exponentially sparse, hindering meaningful learning through standard RL.

In contrast, curriculum learning with adaptive supervision succeeds in solving this task using a limited number of training samples. Suppose the reward threshold is fixed at 0.5. The adaptive curriculum can initially present the input question X along with almost the entire solution—namely Y_1, \ldots, Z_{L-1} —to the model. In this case, the model only needs to generate the final step, Y_L, Z_L . Assuming the model has learned the task format during pretraining, the probability of generating a valid final step is approximately 0.5, in contrast to 2^{-L} when generating the entire sequence.

This setup enables the model to explore and learn the final reasoning step. Specifically, for each partial sequence $X_1,\ldots,X_L,Y_1,\ldots,Z_{L-1}$, the model produces two valid continuations: Y_L,Z_L and Y'_L,Z'_L , where $x'=x\oplus 1$ denotes the complement of x. This variation allows gradient-based learners to infer that Y_L directly influences Z_L : flipping Y_L changes Z_L while all other inputs remain fixed. Consequently, learning Z_L reduces to learning a degree-two parity function, as one input (Y_L) is already learned, which has lower sample complexity than learning the degree-three parity function. In particular, degree-two parities can be learned with $\tilde{\Omega}(L)$ samples (Glasgow, 2024; Kou et al., 2024). Once the model masters the final step, the curriculum is adjusted to reveal a shorter solution prefix (up to Z_{L-2}), requiring the model to generate $Y_{L-1}, Z_{L-1}, Y_L, Z_L$. Since it has already learned to generate Y_L, Z_L , the focus shifts to learning the (L-1)-th step. This process is repeated iteratively by gradually shortening the hint until the model learns the full sequence. This progression demonstrates that there exist sample-size regimes where standard SFT + RL fails, but AdaBack successfully learns the task.

We note that the argument above is simplified, as the actual learning complexity depends on the specific learning model used. We will discuss such subtleties in Appendix B. Here, we focus on the empirical evidence supporting this phenomenon. We consider the case of L=16 with a training set size of n=1024, using the Llama 3.2 1B model (Grattafiori et al., 2024). The reward function assigns a value of 1 to fully correct sequences and 0.1 to sequences that are syntactically valid (i.e., correct number of bits).

We first perform SFT on the training set to help the model learn the task format. We then compare the performance of standard RL and AdaBack on the task. As shown in Figure 2 (left), AdaBack achieves substantial reward early in training and progressively reduces the portion of the solution that is revealed, leading to seamless learning of the task. In contrast, standard RL (Figure 2, right) fails to obtain meaningful rewards and does not learn the task.

This experiment highlights how AdaBack enables guided exploration via partially correct sequences, effectively expanding the class of learnable problems in structured sequence modeling. It also establishes a clear separation between SFT + standard RL and AdaBack: while the former struggles with sparse rewards, AdaBack offers a principled middle ground, avoiding both the limitations of full supervision and the exploration challenges of standard RL.

3 EXPERIMENTS

In this section, we evaluate AdaBack on real-world reasoning tasks. In Section 2.2, we showed that AdaBack enables models to solve problems that are intractable under standard supervised fine-tuning, reinforcement learning, or their combination. Our goal here is to investigate whether AdaBack improves generalization and reasoning ability beyond what standard RL pipelines achieve on real-world reasoning tasks.

We train models using the GRPO algorithm (Shao et al., 2024) on base models of the Llama-3 family (Grattafiori et al., 2024). For fairness of attribution and clarity, we have followed the standard practice of works like (Shao et al., 2024; Xi et al., 2024) by using simple greedy or sampling-based decoding

¹We use \tilde{O} to hide the logarithmic factors.




Figure 2: **Training Dynamics.** *Left:* Training and test rewards along with supervision ratios throughout training. With AdaBack, Llama 3.2 1B successfully learns the task in under 700 iterations. *Right:* Training and test rewards for SFT+RL (red) plateau at 0.1, indicating that only the output format—learned during supervised pretraining—has been retained. Test reward for R3 (Xi et al., 2024) is shown in purple; it reaches only 0.8 reward after more than 16,000 iterations. R3 segments training examples at all whitespace positions and applies RL uniformly over these fragments, resulting in inefficiency due to its non-adaptive strategy.

for our comparisons with a prompt as simple as 'Please provide the final answer to the question in \boxed{}.' All baseline experiments are ran until convergence.

The standard RL recipe involves first performing SFT on rationales, followed by RL fine-tuning (Gulcehre et al., 2023; Guan et al., 2025). However, this paradigm has recently been challenged by models trained solely with RL updates such as GRPO, without any SFT phase (DeepSeek-AI, 2025). Accordingly, we compare four variants: GRPO on a base model, GRPO on an SFT model, AdaBack-GRPO on a base model, and AdaBack-GRPO on an SFT model. Further experimental details are provided in Appendix D.

Generalization on Natural Language Reasoning Tasks We evaluate on two standard mathematical reasoning benchmarks: MATH (Hendrycks et al., 2021) and GSM8k (Cobbe et al., 2021). In line with Mirzadeh et al. (2024) and Li et al. (2024), to assess generalization beyond pretraining exposure, we introduce two new variants of GSM8k:

- In Base-7 GSM8k, all numeric quantities and computations are represented in base-7.2 While the problems themselves are unchanged, this symbolic shift forces the model to reason over a format it has not encountered during pretraining—analogous to deploying a model in a culture with a different numerical base. Performance on this dataset requires the model to generalize beyond familiar surface forms and reasoning chains observed in pretraining.
- We create **Tensor-2 GSM8k** inspired by Hosseini et al. (2024), by concatenating pairs of GSM8k problems and their solutions into single instances, yielding longer reasoning chains. While the resulting chains may be semantically disjoint, this construction increases the number of sequential reasoning steps required in the generation. This setup tests whether models can scale their reasoning over longer inputs and outputs, beyond what is required in the original GSM8k.

Table 1 summarizes performance across these tasks. We find that AdaBack consistently outperforms both GRPO and SFT+GRPO, particularly in out of pretraining distribution settings like Tensor-2 GSM8k. Notably, AdaBack applied directly to the base model often matches or exceeds the performance of the standard RL applied to the base and SFT-initialized models, suggesting that adaptive partial supervision on a base model may serve as a more effective prior than initializing from a SFT checkpoint. Interestingly, AdaBack applied on the base model may even outperform AdaBack applied on the SFT model as in the Tensor-2 experiment. We elaborate on this phenomenon further in the text.

²We dropped (fewer than 1000) samples where the question or answer required division, as base-10 divisions can have periodic representations in base-7.

Table 1: Final test accuracy for each method across tasks and model sizes.

Method	MATH		GSM8k		Base-7 GSM8k		Tensor-2 GSM8k	
	1B	3B	1B	3B	1B	3B	1B	3B
Base+RL	6.4	15.0	7.9	63.7	4.8	4.9	0.0	0.0
SFT+RL	7.4	17.7	36.7	72.7	14.4	45.4	6.9	42.7
AdaBack	9.1	19.1	39.2	73.3	18.4	43.9	8.5	49.2
SFT+AdaBack	9.5	19.9	43.2	70.7	24.5	49.9	11.3	42.2

Per-Sample Curriculum Without Manual Staging Curriculum learning typically requires hand-crafted stages, scheduling heuristics, and careful tuning of hyperparameters—such as how long to train at each difficulty level and with what parameters, and when to transition to harder ones. In contrast, AdaBack performs an automatic per-sample curriculum by adjusting supervision based on a simple reward threshold. Each training point progresses at its own pace, without global stage definitions. Although our implementation uses a stochastic binary search over supervision ratios, this is not essential: linear search or other adaptive strategies could be substituted. The key principle is to allow each example to progress at its own pace, without requiring global curriculum design or extensive hyperparameter tuning. As seen in Figure 3 (left) average portions naturally decrease and rewards increase with AdaBack without any boilerplate curriculum scaffolding. τ is the only hyperparameter needed to be set for AdaBack, and training is not sensitive to it. We have elaborated the choice of τ in Section C.




Figure 3: AdaBack vs. Standard RL Across Model Initializations. Results from training Llama3-1B on GSM8k dataset. The top row shows results for models initialized with SFT, while the bottom row shows base (non-SFT) models. The left column presents AdaBack training dynamics: train reward increases as supervision ratios (portions) decrease. The right column shows standard RL.

Training Dynamics and Initialization Figure 3 shows training curves comparing AdaBack and standard RL across SFT and base model initializations. In both settings, AdaBack achieves better rewards both at training and test compared to standard GRPO training. We provide additional training figure in Appendix F.

Does AdaBack Expand the Model's Solution Space? Yue et al. (2025) argue that RL fine-tuning reweights a model's output distribution without expanding its effective reasoning capacity. To test this claim, we evaluate models using $pass@k^3$, a metric that captures the breadth of plausible solutions. Figure 4 shows that AdaBack significantly improves pass@k over standard RL on both the base and SFT models, especially at large k. These gains are most pronounced when AdaBack is applied to the base model, again suggesting that it facilitates the discovery of new solution modes rather than just refining existing ones. If SFT-free RL with AdaBack introduces novel capabilities, we expect pass@k to increase even when base model coverage is low, which is what we show in Figure 4. In some cases, such as on the Tensor-2 GSM8k (see Table 1), AdaBack on the base model even outperforms AdaBack on the SFT-initialized model, suggesting that SFT may sometimes restrict the search space too early, limiting the benefits of exploration. The latter can be observed in Figure 4 as well.




Figure 4: **Pass@k** for Llama3-1B SFT-initialized models (left) and base models (right) on GSM8k. AdaBack keeps a significant gap compared to standard RL and improves performance at higher k even without SFT suggesting it expands the solution distribution rather than reweighting known answers (contra Yue et al. (2025)).

When Does AdaBack Fail to Help? On the MATH dataset with Llama 3.2 3B-Instruct, we observe no gains from AdaBack. As shown in Figure 8, both train and test rewards saturate quickly, with nearly all questions solved within a few hundred iterations. We observed a similar behavior on other models like the Qwen 2.5 family of models (Yang et al., 2024) as well as in the case of GSM8k for the 3B model in Table 1. This happens when RL reaches perfect training reward within a few iterations during training, i.e., when the dataset is too simple for the model. This suggests that these models have likely encountered much of the dataset during pretraining, leaving little room for RL improvements or for AdaBack to assist exploration. In such cases, AdaBack provides limited benefit—highlighting that its main strength lies in tasks where sparse reward or symbolic mismatch creates a significant learning barrier.

4 RELATED WORK

Here, we provide a summary of the related work. We also provide a more extensive literature review in Appendix A.

Limits of Reinforcement Learning for Reasoning Several studies have highlighted the difficulty of applying reinforcement learning (RL) to structured reasoning tasks. Exploration remains a major bottleneck: reward signals are sparse, and correct reasoning chains are exponentially rare in the output space. Havrilla et al. (2024) empirically evaluated several RL strategies on reasoning benchmarks and found that models fail to discover solutions outside the support of the base SFT model. They tested curriculum-inspired techniques such as *backtracking*—starting the model partway through a solution and gradually shifting the start point earlier Salimans and Chen (2018)—and *prioritized*

 $^{^{3}}$ Pass@k is the probability that at least one out of k model-generated outputs correctly solves a given problem. Chen et al. (2021); Brown et al. (2024)

⁴In an extreme case, the model may just memorize the answers during the SFT phase, removing hope for any exploration.

level replay (PLR)—sampling more frequently from difficult problems Jiang et al. (2020). Despite these interventions, performance gains were negligible: RL primarily amplified answers already assigned non-trivial probability by the pretrained model. These findings align with recent theoretical arguments by Yue et al. (2025), who claim that RL fine-tuning reweighs existing reasoning paths but fails to induce fundamentally new capabilities.

In contrast, we find that curriculum learning with partial supervision can induce such capabilities. Through carefully designed adaptive exposure to partial solutions, our models learn to complete problems that neither SFT nor RL could solve—even partially—highlighting the limitations of standard RL exploration and the promise of fine-grained curricula.

Curriculum Learning Curriculum learning was originally introduced to improve generalization and convergence by training models on simpler examples before harder ones (Bengio et al., 2009). In RL, backtracking (Salimans and Chen, 2018) and RFCL Tao et al. (2024) adopt this principle by initializing rollouts from later points in expert demonstrations and gradually shifting the start point backward.

In language modeling, R3 (Xi et al., 2024) introduces a step-wise curriculum where the model is trained to complete reasoning chains from progressively earlier points. However, despite being motivated as a dynamic curriculum, R3 effectively performs static data augmentation: demonstrations are sliced at delimiter positions (e.g., newlines in GSM8k), and each slice is treated as an independent training sample. In datasets without consistent delimiters, R3 falls back on uniformly partitioning each solution into a fixed number of segments (e.g., five), regardless of problem difficulty or solution length, and then mixing all segments across the dataset. This reliance on global heuristics limits scalability and applicability to domains without clear step boundaries, such as MATH (Hendrycks et al., 2021) or our synthetic parity benchmark. In particular, for the parity task we show that R3 has significantly worse learning speed and generalization in Figure 2. This is consistent with the theoretical results of Cornacchia and Mossel (2023); Abbe et al. (2023b) where it is proven that applying a curriculum over distributions can significantly improve the speed of convergence over simply mixing the distributions.

Our approach differs in both granularity and adaptability. We introduce an adaptive per-sample curriculum that dynamically controls the fraction of the output revealed to the model during training. Unlike R3, we do not rely on external step segmentation or dataset-specific structure. This fine-grained supervision allows the model to incrementally learn reasoning behaviors directly from reward signals, without explicit hinting or hard-coding intermediate targets.

5 CONCLUSION, LIMITATIONS, AND FUTURE WORK

Curriculum learning is often seen as a training heuristic to improve convergence. In this work, we show that per-sample curriculum learning via adaptive partial supervision enables fundamentally new reasoning capabilities that are inaccessible to standard supervised fine-tuning or reinforcement learning alone. Our proposed method, AdaBack, succeeds on a synthetic parity task where both RL and SFT fail, offering a constructive separation result. On real-world datasets such as MATH (Hendrycks et al., 2021), GSM8k (Cobbe et al., 2021), and its variants, AdaBack outperforms baseline RL and SFT+RL pipelines, often matching or exceeding standard RL applied on SFT-initialized models even when applied directly to a base model.

These results suggest that AdaBack not only improves sample efficiency, but also facilitates exploration of new solution modes, as evidenced by consistent gains in pass@k. However, we also observe that AdaBack provides limited benefit when the base model has already memorized much of the task, as appears to be the case with modern reasoning-apt models and math datasets.

A key limitation of our current implementation arises in the large dataset regime. When the number of unique training samples is high, most examples are seen infrequently and therefore rely heavily on global moving averages for supervision scheduling. This may reduce the effectiveness of per-sample adaptation. One promising direction for future work is to define supervision schedules not per sample, but per region in embedding space—e.g., using the average supervision level of the k nearest neighbors of each sample. This would preserve the benefits of adaptation while remaining scalable in data-rich settings.

REFERENCES

486

487

496

497

498

499

504

505

512

513

514

515

519

520

521

522

523 524

525

527

528

529

- E. Abbe and C. Sandon. On the universality of deep learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20061–20072. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/file/e7e8f8e5982b3298c8addedf6811d500-Paper.pdf.
- E. Abbe and C. Sandon. Polynomial-time universality and limitations of deep learning. Communications on Pure and Applied Mathematics, 76(11):3493–3549, 2023. doi: https://doi.org/10.1002/cpa.22121. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.22121.
 - E. Abbe, P. Kamath, E. Malach, C. Sandon, and N. Srebro. On the power of differentiable learning versus PAC and SQ learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- E. Abbe, E. B. Adsera, and T. Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2552–2623. PMLR, 2023a.
 - E. Abbe, E. Cornacchia, and A. Lotfi. Provable advantage of curriculum learning on parity targets with mixed inputs. *Advances in Neural Information Processing Systems*, 36:24291–24321, 2023b.
- E. Abbe, S. Bengio, A. Lotfi, C. Sandon, and O. Saremi. How far can transformers reason? the globality barrier and inductive scratchpad. *Advances in Neural Information Processing Systems*, 37:27850–27895, 2024.
- A. Ahmadian, C. Cremer, M. Gallé, M. Fadaee, J. Kreutzer, O. Pietquin, A. Üstün, and S. Hooker. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms, 2024. URL https://arxiv.org/abs/2402.14740.
 - M. H. Amani, N. M. Baldwin, A. Mansouri, M. Josifoski, M. Peyrard, and R. West. Symbolic autoencoding for self-supervised sequence learning, 2024. URL https://arxiv.org/abs/2402.10575.
- Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, and et. al. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022. URL https://arxiv.org/abs/2204.05862.
 - B. Barak, B. L. Edelman, S. Goel, S. Kakade, E. Malach, and C. Zhang. Hidden progress in deep learning: Sgd learns parities near the computational limit. *arXiv preprint arXiv:2207.08799*, 2022.
 - Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.
 - Y. Bengio, N. Léonard, and A. C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *ArXiv*, abs/1308.3432, 2013.
 - B. Brown, J. Juravsky, R. Ehrlich, R. Clark, Q. V. Le, C. Ré, and A. Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling, 2024. URL https://arxiv.org/abs/2407.21787.
- M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba. Evaluating large language models trained on code, 2021. URL https://arxiv.org/abs/2107.03374.
 - K. Cobbe, C. Hesse, J. Hilton, and J. Schulman. Leveraging procedural generation to benchmark reinforcement learning. *arXiv preprint arXiv:1912.01588*, 2019.

- K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton,
 R. Nakano, C. Hesse, and J. Schulman. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168, 2021.
 - E. Cornacchia and E. Mossel. A mathematical model for curriculum learning for parities. In *International Conference on Machine Learning*, pages 6402–6423. PMLR, 2023.
 - DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.
 - B. Edelman, S. Goel, S. Kakade, E. Malach, and C. Zhang. Pareto frontiers in deep feature learning: Data, compute, width, and luck. *Advances in Neural Information Processing Systems*, 36, 2024.
 - G. Feng, B. Zhang, Y. Gu, H. Ye, D. He, and L. Wang. Towards revealing the mystery behind chain of thought: a theoretical perspective. *Advances in Neural Information Processing Systems*, 36, 2024.
 - C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel. Reverse curriculum generation for reinforcement learning. In *Conference on robot learning*, pages 482–495. PMLR, 2017.
 - S. Gao, A. Bosselut, S. Bengio, and E. Abbe. Abstral: Augmenting llms' reasoning by reinforcing abstract thinking, 2025. URL https://arxiv.org/abs/2506.07751.
 - M. Glasgow. SGD finds then tunes features in two-layer neural networks with near-optimal sample complexity: A case study in the XOR problem. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=HgOJlxzB16.
 - S. Goyal, Z. Ji, A. S. Rawat, A. K. Menon, S. Kumar, and V. Nagarajan. Think before you speak: Training language models with pause tokens. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ph04CRkPdC.
- A. Grattafiori, A. Dubey, and Others. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.
 - X. Guan, L. L. Zhang, Y. Liu, N. Shang, Y. Sun, Y. Zhu, F. Yang, and M. Yang. rstar-math: Small llms can master math reasoning with self-evolved deep thinking, 2025. URL https://arxiv.org/abs/2501.04519.
- C. Gulcehre, T. L. Paine, S. Srinivasan, K. Konyushkova, L. Weerts, A. Sharma, A. Siddhant, A. Ahern, M. Wang, C. Gu, W. Macherey, A. Doucet, O. Firat, and N. de Freitas. Reinforced self-training (rest) for language modeling, 2023. URL https://arxiv.org/abs/2308.08998.
 - A. Havrilla, Y. Du, S. C. Raparthy, C. Nalmpantis, J. Dwivedi-Yu, M. Zhuravinskyi, E. Hambro, S. Sukhbaatar, and R. Raileanu. Teaching large language models to reason with reinforcement learning, 2024. URL https://arxiv.org/abs/2403.04642.
 - D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt. Measuring mathematical problem solving with the math dataset, 2021. URL https://arxiv.org/abs/2103.03874.
- A. Hosseini, A. Sordoni, D. Toyama, A. Courville, and R. Agarwal. Not all llm reasoners are created equal. *arXiv preprint arXiv:2410.01748*, 2024.
 - M. Jiang, E. Grefenstette, and T. Rocktäschel. Prioritized level replay. In *International Conference on Machine Learning*, 2020.
 - L. Kaiser and S. Bengio. Discrete autoencoders for sequence models. ArXiv, abs/1801.09797, 2018.
- M. Kearns. Efficient noise-tolerant learning from statistical queries. J. ACM, 45(6):983–1006, Nov. 1998. ISSN 0004-5411. doi: 10.1145/293347.293351. URL http://doi.acm.org/10.1145/293347.293351.
 - T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.

Y. Kou, Z. Chen, Q. Gu, and S. M. Kakade. Matching the statistical query lower bound for \$k\$-sparse parity problems with sign stochastic gradient descent. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=EbSSBvwUWw.

J. Lanchantin, S. Toshniwal, J. Weston, S. Sukhbaatar, et al. Learning to reason and memorize with self-notes. *Advances in Neural Information Processing Systems*, 36, 2024.

- A. Lewkowycz, A. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *arXiv* preprint arXiv:2206.14858, 2022.
- Q. Li, L. Cui, X. Zhao, L. Kong, and W. Bi. GSM-plus: A comprehensive benchmark for evaluating the robustness of LLMs as mathematical problem solvers. In L.-W. Ku, A. Martins, and V. Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2961–2984, Bangkok, Thailand, Aug. 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.163. URL https://aclanthology.org/2024.acl-long.163/.
- B. Liu, S. Bubeck, R. Eldan, J. Kulkarni, Y. Li, A. Nguyen, R. Ward, and Y. Zhang. Tinygsm: achieving >80 URL https://arxiv.org/abs/2312.09241.
- W. Merrill and A. Sabharwal. The expressive power of transformers with chain of thought. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=NjNGlPh8Wh.
- I. Mirzadeh, K. Alizadeh, H. Shahrokhi, O. Tuzel, S. Bengio, and M. Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models, 2024. URL https://arxiv.org/abs/2410.05229.
- S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone. Curriculum learning for reinforcement learning domains: A framework and survey. *The Journal of Machine Learning Research*, 21(1):7382–7431, 2020.
- M. Nye, A. J. Andreassen, G. Gur-Ari, H. Michalewski, J. Austin, D. Bieber, D. Dohan, A. Lewkowycz, M. Bosma, D. Luan, C. Sutton, and A. Odena. Show your work: Scratchpads for intermediate computation with language models, 2021.
- T. Olausson, A. Gu, B. Lipkin, C. Zhang, A. Solar-Lezama, J. Tenenbaum, and R. Levy. LINC: A neurosymbolic approach for logical reasoning by combining language models with first-order logic provers. In H. Bouamor, J. Pino, and K. Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5153–5176, Singapore, Dec. 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.313. URL https://aclanthology.org/2023.emnlp-main.313/.
- R. Y. Pang, W. Yuan, K. Cho, H. He, S. Sukhbaatar, and J. Weston. Iterative reasoning preference optimization, 2024a. URL https://arxiv.org/abs/2404.19733.
- R. Y. Pang, W. Yuan, K. Cho, H. He, S. Sukhbaatar, and J. E. Weston. Iterative reasoning preference optimization. *ArXiv*, abs/2404.19733, 2024b. URL https://api.semanticscholar.org/CorpusID:269457506.
- N. Razin, H. Zhou, O. Saremi, V. Thilak, A. Bradley, P. Nakkiran, J. Susskind, and E. Littwin. Vanishing gradients in reinforcement finetuning of language models, 2024. URL https://arxiv.org/abs/2310.20703.
- C. Resnick, R. Raileanu, S. Kapoor, A. Peysakhovich, K. Cho, and J. Bruna. Backplay: "man muss immer umkehren", 2022. URL https://arxiv.org/abs/1807.06919.
 - T. Salimans and R. J. Chen. Learning montezuma's revenge from a single demonstration. *ArXiv*, abs/1812.03381, 2018.

- R. J. Sánchez, L. Conrads, P. Welke, K. Cvejoski, and C. O. Marin. Hidden schema networks. In
 A. Rogers, J. L. Boyd-Graber, and N. Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto,
 Canada, July 9-14, 2023, pages 4764–4798. Association for Computational Linguistics, 2023.
 - D. Saxton, E. Grefenstette, F. Hill, and P. Kohli. Analysing mathematical reasoning abilities of neural models. *arXiv preprint arXiv:1904.01557*, 2019.
 - T. Schaul, J. Quan, I. Antonoglou, and D. Silver. Prioritized experience replay. *CoRR*, abs/1511.05952, 2015. URL https://api.semanticscholar.org/CorpusID:13022595.
 - J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
 - Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. K. Li, Y. Wu, and D. Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.
 - S. Sukhbaatar, E. Denton, A. Szlam, and R. Fergus. Learning goal embeddings via self-play for hierarchical reinforcement learning, 2018a. URL https://arxiv.org/abs/1811.09083.
 - S. Sukhbaatar, Z. Lin, I. Kostrikov, G. Synnaeve, A. Szlam, and R. Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play, 2018b. URL https://arxiv.org/abs/1703.05407.
 - S. Tao, A. Shukla, T.-k. Chan, and H. Su. Reverse forward curriculum learning for extreme sample and demonstration efficiency in rl, 2024.
 - J. Uesato, N. Kushman, R. Kumar, F. Song, N. Siegel, L. Wang, A. Creswell, G. Irving, and I. Higgins. Solving math word problems with process- and outcome-based feedback, 2022. URL https://arxiv.org/abs/2211.14275.
 - A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
 - P. Veličković, A. P. Badia, D. Budden, R. Pascanu, A. Banino, M. Dashevskiy, R. Hadsell, and C. Blundell. The clrs algorithmic reasoning benchmark. *arXiv* preprint arXiv:2205.15659, 2022.
 - J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
 - R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.
 - Z. Xi, W. Chen, B. Hong, S. Jin, R. Zheng, W. He, Y. Ding, S. Liu, X. Guo, J. Wang, H. Guo, W. Shen, X. Fan, Y. Zhou, S. Dou, X. Wang, X. Zhang, P. Sun, T. Gui, Q. Zhang, and X. Huang. Training large language models for reasoning through reverse curriculum reinforcement learning, 2024.
 - S. Xue, Z. Huang, J.-Y. Liu, X. Lin, Y. Ning, B. Jin, X. Li, and Q. Liu. Decompose, analyze and rethink: Solving intricate problems with human-like reasoning cycle. In *Neural Information Processing Systems*, 2024. URL https://api.semanticscholar.org/CorpusID: 276117098.
 - A. Yang, B. Yang, B. Hui, B. Zheng, B. Yu, C. Zhou, C. Li, C. Li, D. Liu, F. Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
 - K. Yang, A. Swope, A. Gu, R. Chalamala, P. Song, S. Yu, S. Godil, R. J. Prenger, and A. Anandkumar. Leandojo: Theorem proving with retrieval-augmented language models. *Advances in Neural Information Processing Systems*, 36:21573–21612, 2023.
 - Y. Yue, Z. Chen, R. Lu, A. Zhao, Z. Wang, Y. Yue, S. Song, and G. Huang. Does reinforcement learning really incentivize reasoning capacity in Ilms beyond the base model?, 2025. URL https://arxiv.org/abs/2504.13837.

E. Zelikman, Y. Wu, J. Mu, and N. Goodman. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.

E. Zelikman, G. Harik, Y. Shao, V. Jayasiri, N. Haber, and N. D. Goodman. Quiet-star: Language models can teach themselves to think before speaking, 2024. URL https://arxiv.org/abs/2403.09629.

A EXTENDED RELATED WORK

Scratchpads and Chain-of-Thought Strategies It is widely believed that success on challenging problems requires a model to know how to use intermediary computations in the context, to reason and deduct the final answers. Nye et al. (2021) proposed supervised training of Transformers to use scratchpads in addition to final answer, showing improvements on tasks like executing Python code and evaluating polynomials. Similarly, Wei et al. (2023) proposed chain-of-thought prompting, showing that large language models can generate scratchpads via in-context demos and without explicit training. Moreover, Kojima et al. (2022) studied zero-shot chain-of-thought generation for language models. Lanchantin et al. (2024) introduced the concept of self-notes, showing benefits of interleaving the intermediate reasoning steps within the question/context. Goyal et al. (2024) introduced pause tokens which act as place-holder tokens providing models with more computation time before output generations.

Several works have shown that allowing transformers to produce a chain-of-thought would increase their expressivity (Feng et al., 2024; Merrill and Sabharwal, 2024). Further, Abbe et al. (2024) put forward the notion of globality degree of a task as a hardness measure and show that scratchpads can make learning more efficient by breaking the globality of a task. They also proposed inductive scratchpads which impose a Markovian structure over reasoning steps, improving length generalization. Gao et al. (2025) proposed AbstRaL, showing improved robustness by training language models on mathematical abstraction of reasoning traces, instead of natural language chain-of-thoughts. Outside of reasoning and natural language scratch pads, unsupervised learning of intermediary symbolic sequences using straight-through gradient estimators (Bengio et al., 2013) has been studied in (Amani et al., 2024; Sánchez et al., 2023; Kaiser and Bengio, 2018).

RL and Reasoning Benchmarks Reinforcement learning has emerged as a dominant paradigm for post-training language models on tasks involving sparse, verifiable reasoning signals Havrilla et al. (2024). Our reasoning tasks and settings are closest to STaR (Zelikman et al., 2022), Quiet-STaR (Zelikman et al., 2024), and R-STaR-MATH (Guan et al., 2025), which apply RL to improve mathematical and symbolic reasoning through selective training on verified rationales.

In parallel, alignment with human preferences has also leveraged RL techniques—most prominently Proximal Policy Optimization (PPO (Schulman et al., 2017)) (Bai et al., 2022; Pang et al., 2024a). Recent efforts have explored variants of REINFORCE-style updates and reward model bootstrapping, often using iterative fine-tuning pipelines. Interestingly, multiple works have reported that iterative filtering and fine-tuning on correct completions can match or exceed PPO performance in some domains (Gulcehre et al., 2023; Ahmadian et al., 2024).

Two standard benchmarks for mathematical reasoning used in most of the just-mentioned work in RL for reasoning are GSM8k (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) datasets. GSM8k consists of grade-school arithmetic problems requiring multi-step solutions, often accompanied by natural language justifications. MATH contains higher-difficulty competition-style math problems with structured step-by-step solutions. Nevertheless, RL approaches can be applied to any task for which verifiers or other forms of reward functions are available, including coding (Chen et al., 2021) and formal math (Yang et al., 2023).

Alternatively to RL, some natural language questions can be translated into formal logic and answered using automated theorem provers—for instance, Olausson et al. (2023) maps questions into first-order logic using an LLM, then delegates inference to a symbolic prover. AbstRaL also solves GSM8k-style problems by first converting them into abstract symbolic equations and then using a solver (Gao et al., 2025).

Process vs Outcome Supervision Uesato et al. (2022) compared process-based feedback (reward each correct intermediate step) with outcome-based feedback (reward only at the final answer) on GSM8k. They found final answer accuracy can be similar with outcome-only training, but the quality of the reasoning steps was much higher with process supervision—the process supervision reducing reasoning errors dramatically. This could be another motivation for backtracking methods as the model learns by reinforcing small steps conditioned on correct chain-of-thoughts.

Curriculum Learning in Reinforcement Learning Several other works in reinforcement learning have explored curriculum strategies to overcome sparse reward challenges (Narvekar et al., 2020). Florensa et al. (2017) uses a reverse curriculum for training, where starting states become increasingly difficult during training. However, compared to Salimans and Chen (2018), these starting states do not come from demonstrations. Backplay (Resnick et al., 2022) demonstrated strong gains in environments like Pommerman and Atari by starting episodes near the goal state using a single demonstration and gradually moving the starting point backward, enabling the agent to outperform the suboptimal demonstrator. Prioritized Level Replay (PLR) (Jiang et al., 2020) focuses training on the hardest levels in procedural environments like Procgen(Cobbe et al., 2019) by adaptively replaying levels where the agent performs poorly. Sukhbaatar et al. (2018b) proposed an automatic curriculum where a teacher agent proposes increasingly difficult tasks for a learner agent, leading to emergent complex curriculum strategies without any hand-designed task progression. This was extended in Sukhbaatar et al. (2018a) to learn goal embeddings and reusable low-level policies through self-play.

Comparison with State-of-the-art Mathematical Reasoning Models Although recent leader-board results show strong performance, they are not directly comparable to our work because they typically rely on much larger models, additional math-heavy pretraining datasets, or complex prompting and inference-time searching strategies Xue et al. (2024); Pang et al. (2024b); Liu et al. (2023). In contrast, our focus is on isolating the contribution of AdaBack itself by comparing it fairly against standard RL baselines under controlled settings. Importantly, AdaBack is not an alternative to these methods but a complementary training enhancement. It can be combined with larger models, post-training strategies, or prompting techniques like chain-of-thought prompting and majority voting Xue et al. (2024) without overhead. This makes AdaBack a foundational component that strengthens models from the ground up and can integrate seamlessly into more complex state-of-the-art pipelines Pang et al. (2024b).

B DISCUSSION ON LEARNING PARITIES

The problem of learning parity functions has been extensively studied in the theory of machine learning (Abbe and Sandon, 2020; 2023; Barak et al., 2022; Abbe et al., 2023a; Edelman et al., 2024; Glasgow, 2024). Typically, the task is defined as follows⁵: each bit in the input is sampled independently and uniformly from the Rademacher distribution $\operatorname{Rad}(1/2)$ (i.e., -1 or 1 with equal probability). The target function is a parity over a subset of bits, i.e., $\prod_{i \in S} x_i$ for some $S \subseteq [d]$. The size of S, denoted |S|, is called the degree of the parity. When $|S| = O_d(1)$ is constant (with respect to the input dimension), the problem is referred to as learning a sparse parity.

Due to symmetry, we often discuss learning a parity of degree k without specifying the particular subset S. Most learning formulations assume the degree k is known while the identity of S must be learned, yielding a hypothesis class of size $\binom{d}{k}$. It is well established that the difficulty of learning sparse parities increases with the degree k. In the statistical query (SQ) model (Kearns, 1998), it has been shown that learning a parity of degree k requires $\Omega(d^k)$ queries, which translates to $\Omega(d^{k-1})$ samples (Kou et al., 2024).

Similar lower bounds have been conjectured/shown for neural networks. In particular, fully connected networks with bounded width and depth and rotationally invariant weight initializations are conjectured to require $\Omega(d^{\max(k-1,1)})$ online gradient steps to learn a degree-k parity (Abbe et al., 2023a). This has been proven under specific assumptions, such as for noisy population gradient descent (Abbe and Sandon, 2023; Abbe et al., 2021). Recently, Kou et al. (2024) showed that this lower bound is tight: a variant of SGD (namely, sign-SGD) can learn the task in $O(\log d)$ iterations using batches of size $\tilde{O}(d^{k-1})$. The problem becomes more subtle in the offline setting where samples can be reused. In this case, Abbe et al. (2023a) conjecture the optimal sample complexity to be $\Theta(d^{\max(k/2,1)})$. We note that one can often reduce the sample complexity at the cost of using larger models, where the width/depth of the model scales with d (Edelman et al., 2024).

 $^{^{5}}$ We follow standard theoretical conventions and use ± 1 values for the bits. Equivalent results hold for 0,1 bits. Moreover, since Transformers embed inputs into continuous vectors, empirical results are invariant to the bit representation.

In our proposed chain-of-parities task, an input sequence of length L contains L-1 parity targets of degree 3 and one initial parity of degree 2. Learning all of these parity targets jointly with a single model may be harder (due to interference) or easier (due to shared information, such as the influence of previous bit), depending on the model architecture—particularly positional embeddings in Transformers. A simplifying assumption is to treat the learning of each parity target as independent. Under this assumption, the problem reduces to learning parities of degree 3, for which, based on the discussions above, one would not expect learnability using only $\tilde{\Theta}(L)$ samples. In contrast, the degree-2 parity should be learnable with $\tilde{\Theta}(L)$ samples (Glasgow, 2024; Kou et al., 2024).

This explains why supervised fine-tuning with $\tilde{\Theta}(L)$ samples fails to weakly learn the degree-3 parity targets Z_i for i>1. Consequently, SFT followed by standard reinforcement learning also fails, as the RL phase remains unguided which makes the probability of stumbling upon a valid completion (and receiving reward) exponentially small.

We now explain why AdaBack can succeed where SFT and RL fail. Consider a scenario in which the first L-1 steps of the solution, i.e., $Y_1, Z_1, \ldots, Y_{L-1}, Z_{L-1}$, are provided, and the model must complete Y_L, Z_L . In this setup, there are two valid completions: (Y_L, Z_L) and its complement (Y'_L, Z'_L) , where x' denotes the bitwise complement of x. The fact that the same prefix $X_1, \ldots, X_L, Y_1, \ldots, Z_{L-1}$ admits two such completions indicates that if Z_L is a parity of the previous coordinates, then Y_L lies in the support of this parity. This effectively reduces the problem of learning a degree-3 parity function for Z_L to that of learning a degree-2 parity function (since one bit is already revealed)—a task that is learnable with $\tilde{\Theta}(L)$ samples. This introduces a setting where AdaBack is able to learn the task, while the combination of SFT and standard RL fails.

We do not explore the exact training dynamics here, as they depend on the model architecture and optimization details. Instead, we empirically validate this separation result in Sec. 2.2.

C METHOD DETAILS

Improving Convergence via $\rho=0$ Injection We observed that always sampling supervision ratios from $\rho_t^{(i)} \sim \mathcal{U}(\rho_{\min}^{(i)}, \rho_{\max}^{(i)})$ led to slower convergence and a distribution mismatch between training and validation. Specifically, when a large portion of samples remain difficult (i.e., $\rho_{\min}^{(i)}>c$ for many i), the model would be significantly less exposed to fully unsupervised completions during training, while validation is always conducted at $\rho=0$.

To address this, we introduced a small amount of supervision-free training into the curriculum: with 10% probability, we sample $\rho_t^{(i)}=0$ directly, while using the uniform interval sampling with the remaining 90% probability. This stochastic exposure to 0-portion training helped close the train-test gap and accelerated convergence in the early phases of AdaBack training.

Bootstrapping via Global Averages For samples with no reward history, we initialize $\rho^{(i)}$ from global moving averages $\bar{\rho}_{\min}$ and $\bar{\rho}_{\max}$, which are updated over time using exponential moving averages:

$$\bar{\rho}_{\min} \leftarrow \alpha \rho_{\min}^{(i)} + (1 - \alpha) \cdot \bar{\rho}_{\min}$$
$$\bar{\rho}_{\max} \leftarrow \alpha \rho_{\max}^{(i)} + (1 - \alpha) \cdot \bar{\rho}_{\max}$$

Here $\alpha \in [0,1]$ is a momentum parameter controlling the smoothing rate.

Choice of the reward threshold τ Curriculum learning often requires careful hyperparameter tuning—such as deciding how long to train at each stage and with what optimizer parameters or learning rate. AdaBack simplifies this process by collapsing these decisions into a single parameter: the reward threshold τ . Intuitively, τ controls the average reward the model receives; the binary search mechanism ensures that the supervision is adapted to maintain this target. We experimented with several values of τ and found the final results to be largely robust, except near extreme values (i.e., $\tau \approx 0$ or $\tau \approx 1$). We chose $\tau = 0.5$ in our main experiments, as this maximizes reward variance, following the insights of Razin et al. (2024) on avoiding vanishing gradients for PPO updates.

D EXPERIMENT DETAILS

Computation and Budget We conducted experiments using small language models ranging from 1B to 3B parameters. Input sequence lengths varied across tasks, with a maximum input length of 2048 tokens and generation lengths up to 2048 tokens. Depending on model size and sequence length, we used nodes equipped with either 4 or 8 NVIDIA A100 GPUs. Including all diagnostic and development runs, our total compute usage amounted to approximately 80,000 A100 GPU hours.

Reinforcement Learning Setup We used the GRPO algorithm for all RL experiments, with 8 rollouts (generations) per input sample. Unless otherwise noted, we trained all models for at least 10,000 iterations to ensure convergence to long-term behavior. To further ensure the fairness of comparisons, we continued each standard GRPO training at least as long as the corresponding AdaBack training. Exploring different learning rates, we found $lr=10^{-6}$ to give us stable and rather fast training. We report final test accuracy or reward as the average over the last 5 checkpoints. In rare cases where performance deteriorated at the end of training, we averaged over the last 5 iterations with increasing reward.

Our RL batch size was 256. We did not use a KL penalty or entropy regularization, neither in the loss nor in the reward.

For supervised-fine tuning of the base models, we used a validation set to adjust the hyperparameters and also the number of fine tuning iterations.

For datasets like Tensor-2 GSM8k and the chain-of-parities task, where the output format is nontrivial, we applied a format reward of $r_{\rm format}=0.1$ to encourage structured outputs. This was unnecessary for models initialized from SFT, which already generate syntactically valid sequences and thus begin training with an initial reward near 0.1. However, for consistency, we kept this term in all settings.

Chain-of-Parities Task Setup We use 1024 labeled examples with sequence length L=16. In this setting, the probability of generating a valid solution by chance is 2^{-16} , making reward signals extremely sparse. As a result, standard RL fails to make progress, even when initialized from an SFT checkpoint. The curriculum introduced by AdaBack mitigates this by incrementally revealing intermediate reasoning steps, allowing learning to proceed from the final step backward.

E LLM USAGE

We have used ChatGPT and Gemini models to improve text quality at phrase and paragraph level.

F ADDITIONAL FIGURES

We show how portions and the train and test rewards evolve for GSM8k, MATH, Base-7 GSM8k, and Tensor-2 GSM8k in Figures 3, 5, 6, and 7 respectively.




Figure 5: AdaBack on MATH. Results from training Llama3-3B base model (non-SFT) on MATH dataset. The left column presents AdaBack training dynamics and the right column shows standard RL for comparison.




Figure 6: **AdaBack on Base-7 GSM8k.** Results from training Llama3-1B fine-tuned model (SFT) on Base-7 GSM8k dataset. The left column presents AdaBack training dynamics and the right column shows standard RL for comparison.




Figure 7: AdaBack on Tensor-2 GSM8k. Results from training Llama3-1B fine-tuned model (SFT) on Tensor-2 GSM8k dataset. The left column presents AdaBack training dynamics and the right column shows standard RL for comparison. Note that for this task, outputting an answer with the correct format has 0.1 reward. So the observed rewards are higher than the actual accuracies.




Figure 8: Lack of Learning Signal on MATH. On the MATH dataset, Llama 3.2 3B-Instruct shows minimal learning dynamics, with both train and test rewards saturating early. Nearly all questions are solved within a few hundred iterations, leaving no room for AdaBack to provide further benefit. This may be due to the model having been heavily exposed to examples similar to this dataset during pretraining.